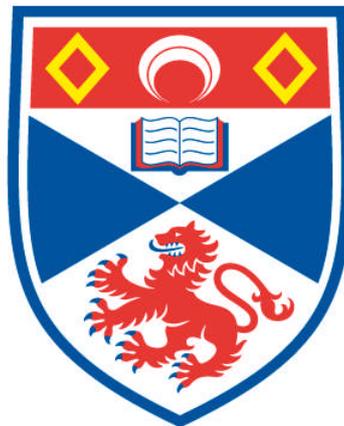


THE SENSORIMOTOR THEORY OF PERCEPTUAL EXPERIENCE

David Silverman

**A Thesis Submitted for the Degree of PhD
at the
University of St Andrews**



2014

**Full metadata for this item is available in
Research@StAndrews:FullText
at:**

<http://research-repository.st-andrews.ac.uk/>

Please use this identifier to cite or link to this item:

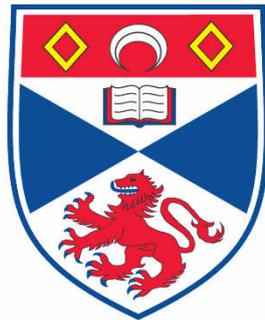
<http://hdl.handle.net/10023/5544>

This item is protected by original copyright

**This item is licensed under a
Creative Commons Licence**

The Sensorimotor Theory of Perceptual Experience

David Silverman



This thesis is submitted in partial fulfilment for the degree of PhD at the University of St Andrews and University of Stirling Graduate Programme in Philosophy (SASP)

30th June 2014

Abstract

The sensorimotor theory is an influential, non-mainstream account of perception and perceptual consciousness intended to improve in various ways on orthodox theories. It is often taken to be a variety of enactivism, and in common with enactivist cognitive science more generally, it de-emphasises the theoretical role played by internal representation and other purely neural processes, giving theoretical pride of place instead to interactive engagements between the brain, non-neural body and outside environment.

In addition to offering a distinctive account of the processing that underlies perceptual consciousness, the sensorimotor theory aims to offer a new and improved account the logical and phenomenological character of perceptual experience, and the relation between physical and phenomenal states. Since its inception in a 2001 paper by O'Regan and Noë, the theory has prompted a good deal of increasingly prominent theoretical and practical work in cognitive science, as well as a large body of secondary literature in philosophy of cognitive science and philosophy of perception. In spite of its influential character, many of the theory's most basic tenets are incompletely or ambiguously defined, and it has attracted a number of prominent objections.

This thesis aims to clarify the conceptual foundations of the sensorimotor theory, including the key theoretical concepts of sensorimotor contingency, sensorimotor mastery, and presence-as-access, and defends a particular understanding of the respective theoretical roles of internal representation and behavioural capacities. In so doing, the thesis aims to highlight the sensorimotor theory's virtues and defend it from some leading criticisms, with particular attention to a response by Clark which claims that perception and perceptual experience plausibly depend on the activation of representations which are not intimately involved in bodily engagements between the agent and environment. A final part of the thesis offers a sensorimotor account of the experience of temporally extended events, and shows how with reference to this we can better understand object experience.

1. Candidate's declarations:

I, David Silverman, hereby certify that this thesis, which is approximately 75,000 words in length, has been written by me, and that it is the record of work carried out by me, or principally by myself in collaboration with others as acknowledged, and that it has not been submitted in any previous application for a higher degree.

I was admitted as a research student in September 2012 and as a candidate for the degree of PhD in September 2010; the higher study for which this is a record was carried out in the University of St Andrews between 2010 and 2014.

Date signature of candidate

2. Supervisor's declaration:

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree.

Date signature of supervisor

3. Permission for publication: (to be signed by both candidate and supervisor)

In submitting this thesis to the University of St Andrews I understand that I am giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. I also understand that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker, that my thesis will be electronically accessible for personal or research use unless exempt by award of an embargo as requested below, and that the library has the right to migrate my thesis into new electronic forms as required to ensure continued access to the thesis. I have obtained any third-party copyright permissions that may be required in order to allow such access and migration, or have requested the appropriate embargo below.

The following is an agreed request by candidate and supervisor regarding the publication of this thesis:

PRINTED COPY

- a) No embargo on print copy

ELECTRONIC COPY

- a) No embargo on electronic copy

Date signature of candidate signature of supervisor.....

Contents

Acknowledgements.....	v
1 - Introduction: The Sensorimotor Theory of Perceptual Experience.....	1
1.1 - 4E Cognitive Science and the Varieties of Enactivism.....	3
1.2 - Structure of Thesis.....	6
2 - Making Conceptual Space.....	8
2.1 - Kepler and Descartes: The Lessons from History.....	8
2.1.1 - Pictures, Filling-In and the Bridge Locus.....	10
2.1.2 - Neural Representation.....	16
2.1.3 - Varieties of Homuncularity.....	19
2.2 - Conservative and Radical Strands.....	23
2.2.1 - The Tenets of Sensorimotor Theory: A Brief Outline.....	25
3 - The Tenets of Sensorimotor Theory.....	31
3.1 - The World as an Outside Memory.....	31
3.2 - Phenomenal Qualities: The Sensorimotor Contingency Account.....	37
3.2.1 - The "Qualia" Problem.....	38
3.2.2 - Qualia and Neural-Correlates of Consciousness.....	39
3.2.3 - The Sensorimotor Contingency Account.....	41
3.2.4 - SMCs as Counterfactual Conditionals.....	44
3.2.5 - Perspectival Properties and Objective Properties.....	47
3.2.6 - Bodiliness, Insubordinateness, Richness and Grabbiness.....	51
3.3 - Sensorimotor Mastery.....	53
3.3.1 - Sensorimotor Attunement.....	54
3.3.2 - Sensorimotor Tracking.....	62
3.3.3 - Sensorimotor Understanding.....	65
3.4 - Perceptual Presence: As-Access and In-Absence.....	67
3.5 - Perception and Dynamical Entanglement.....	77
3.6 - The Extensive Conscious Mind.....	81

4 - Mastering Sensorimotor Understanding.....	93
4.1 - 'Action in Perception' in Action in Perception.....	93
4.2 - Which Direction of Fit?.....	99
4.3 - Sensorimotor Understanding as Subpersonal Representation.....	104
4.4 - Sensorimotor Understanding in Bodily Capacities.....	109
5 - Phenomenal Character: A Sensorimotor-Representation Hybrid.....	116
5.1 - The Non-Sensorimotor, Skill-Based Approach.....	116
5.2 - The Policy of Containment: A Conservative Response.....	122
5.3 - A Challenge to the Policy of Containment.....	124
5.3.1 - The WAY problem.....	125
5.3.2 - Sense and Sensorimotor Contingencies.....	128
5.4 - Mary and Inverted Earth: A Sensorimotor Response.....	134
5.5 - Relation to Rupert's 'Massively Representational Mind'.....	139
5.6 - Chapter Conclusion.....	143
6 - Sorting, Sifting and Other Sensorimotor Skills.....	144
6.1 - Avoiding Zombie Intuitions.....	144
6.2 - Sorting and Sifting as Sensorimotor Skills.....	149
6.3 - Homunculi that Sort and Sift.....	153
6.4 - Radical Anti-Representationalism.....	157
6.5 - Representation Restricted.....	163
7 - Object Experience as Temporal Experience.....	170
7.1 - The Puzzle of Temporal Experience.....	170
7.2 - Coupling and Knowing in Temporal Experience.....	174
7.3 - Object Experience and Skill-Based Access.....	176
7.4 - Extensionalism About Event and Object Experience.....	178
7.5 - Chapter Conclusion.....	181
8 - Conclusion.....	183
Bibliography.....	189

Acknowledgements

First and foremost I offer my sincere gratitude to Professor Michael Wheeler for providing regular and extensive feedback on my work and general support throughout nearly four years of PhD study. I could not have hoped for a better primary supervisor and feel enormously lucky to have had the benefit not only of his vast expertise but also his fantastically conscientious and supportive approach to supervision.

I am grateful to my secondary supervisor, Dr Simon Prosser, for his efforts in looking over draft chapters of this thesis and feedback which has provided a different and useful perspective. Thanks are also due to the other faculty members who have been kind enough to provide comments on my written work. These include Dr Zoe Drayson, who has provided helpful feedback both during and outside of my annual review, and Dr Adrian Haddock, who led two of my annual reviews and has been valuable in encouraging me to more clearly articulate my understanding of the sensorimotor theory.

I would like to thank my fellow students and all the staff members at the St Andrews and Stirling philosophy graduate programme, particularly those based at the philosophy department in Stirling, for providing a friendly close-knit environment and exposure to so many different ways of doing philosophy.

Last but not least I thank my parents, Erika and Julian Silverman. In particular, I am grateful to my father for setting off my interest in philosophy of mind and consciousness after introducing me to the work of Searle, and both for providing moral and often financial support during my almost never-ending years as a student.

1 - The Sensorimotor Theory of Perceptual Experience

The sensorimotor theory is an influential, non-mainstream account of perception and perceptual consciousness intended to offer an improvement on orthodox theories. It is often taken to be a variety of ‘enactivism’, and in common with enactivist cognitive science more generally, it de-emphasises the theoretical role played by internal representation and other purely neural processes, giving theoretical pride of place instead to interactive engagements between the brain, non-neural body and outside environment.

Although orthodox cognitive scientific accounts of vision vary in their details, they commonly descend from the hugely influential work of Marr, brought together in his posthumously published book *Vision* (1982). Marr’s contention was that vision is a form of information processing, beginning with the retinal image and resulting in a detailed, neurally-realised representation of the outside environment, the deployment of which is sufficient for visual awareness. Allied to this general approach is a naive and, as we will see, implausible conception of vision, in which this representation takes the form of an expansive and richly-detailed topographic model of the outside world, the activation of which is sufficient for a richly-detailed visual experience.

Although orthodox theories never explicitly endorse the naive conception, it is not disowned by Marr’s basic claim that vision is a process of constructing a detailed internal representation based on the more limited information available from the retinal image. This has prompted sensorimotor theorists to suggest that research on vision is in many cases tacitly motivated by the naive conception, causing it to generate faulty results. Part of the sensorimotor theory’s ambition, in offering an alternative to Marr’s basic thesis, is to provide a theoretical framework which avoids conceptual mistakes and enables more fruitful empirical research. To this end, the theory claims that perception is not a process of representation, but of skill-driven bodily exploration of the outside environment.

In addition to offering a distinctive account of the processing that underlies perceptual consciousness, the sensorimotor theory aims to explain perceptual consciousness itself, in particular its phenomenal character and the relation between phenomenal and physical properties. Many people continue to be moved by anti-physicalist intuitions about

‘qualia’, and the sensorimotor theory aims to respond to these by providing compelling and principled reasons for thinking that there is nothing more to being in a given phenomenal state than being engaged in a particular kind of embodied interaction with the environment. As before, the theory is supposed to be an advance on accounts of consciousness that appeal solely to facts about the brain, which sensorimotor theorists think lack the resources to defeat anti-physicalist intuitions.

Besides responding to this issue, the sensorimotor theory provides a number of distinctive insights into the logical and phenomenological character of perceptual experience. The theory, interestingly, integrates its contribution to various non-empirically oriented debates about perceptual experience with its account of the cognitive science of perception. In addition to taking a critical stance toward internal representation in cognitive science, the theory rejects representationalism in a number of other domains, including the attempt to explain phenomenal qualities by appeal to representational content, and representational realist accounts of the epistemology and metaphysics of perception. In place of representation, the theory appeals in each case to a bodily relation the perceiver stands in to the outside environment.

The theory was given its best known early statement in a paper by O’Regan and Noë, titled ‘A Sensorimotor Account of Vision and Visual Consciousness’ (2001a) - in fact, it is an account of perception more generally, although the focus is on vision. I will henceforth use the abbreviation O&N to refer to that paper or the paper’s authors in cases where it is that particular paper I am referring to. The other best known statement of the sensorimotor theory is Noë’s (2004) book *Action in Perception*, which is joined by more recent book-length accounts by O’Regan (2011) and Noë (2012), and a number of early papers on the theory, frequently authored in collaboration with Myin (e.g., Myin and O’Regan, 2002; O’Regan, Myin and Noë, 2004; 2005). Although O&N is the first official statement of the theory, it bears strong similarities to an earlier account by Hurley (1998). O&N were not informed by it when they first formulated the account, but Noë went on to collaborate with Hurley (Hurley and Noë, 2003; 2007), suggesting a case can be made for considering Hurley’s book the earliest work in the canon. Other important early work by the theory’s originators, although non-canonical, includes O’Regan (1992), Pessoa, Thompson and Noë (1998), and Thompson

and Noë (2004). The above accounts, and others by the theory's instigators, are supplemented by a large body of secondary work on the theory.

The theory has had an impact in a range of sometimes only loosely-related disciplines. In particular, it has inspired a good deal of increasingly prominent theoretical and practical work in cognitive science: widely-discussed glosses on the theory from the very recent past include, for example, a dynamical systems model by Burhmann, Di Paolo and Barandiaran (2013) and a predictive coding account by Seth (2014). The theory has attracted interest and no little controversy among philosophers of cognitive science, prompting prominent lines of criticism, for instance, by Clark (Clark and Toribio, 2001; Clark, 2006; 2008; 2009), Block (2001; 2005; O'Regan and Block, 2012), Dennett (2001), Prinz (2006), and many others. It has also attracted attention from philosophers of perception, a prominent example of which is a published symposium featuring responses to the sensorimotor theory by Campbell (2008), Martin (2008) and Kelly (2008).

1.1 - 4E Cognitive Science and the Varieties of Enactivism

The sensorimotor theory is situated amongst a complex web of theoretical ancestors and close cousins, and I will draw attention to some of these later in the thesis, in cases where it is helpful to do so. For now, it will be useful to do just a little basic scene setting. Although the theory has had an impact in non-empirically oriented debates about mind and perception, it is primarily a philosophically-motivated entry into the cognitive science of perception and perceptual consciousness. In particular, it enters the running alongside other 4E accounts, a broad class of approaches that depart from the view, mainstream in recent decades, that cognition is best explained by nothing more than the brain's construction and deployment of internal representations.

4E cognitive science claims that cognition is properly characterised using one or more E-words: embedded, embodied, extended and/or enactive. Embedded cognition is cognition that must be explained by appeal to non-trivial facts about the environment in which the cognitive tasks are carried out. Embodied cognition is cognition whose architecture

incorporates the non-neural body as well as the brain.¹ Extended cognition is cognition whose architecture also incorporates the external environment. Enactive accounts frequently reject internal representation altogether, and claim that cognition is ‘enacted’ through a process of interactive engagement between the brain, non-neural body and outside environment. It is worth noting that the degree to which the competing 4E approaches can be reconciled or, alternatively, constitute quite separate and incompatible traditions, is a matter of debate: see, e.g., Ward and Stapleton (2012) for an argument that they can be brought together, and Chemero (2009) and Wheeler (2014) for arguments to the opposite effect.

As I have mentioned, the sensorimotor theory is often regarded as a kind of enactivism. The best-established version of this approach is the ‘biological’ variant, notable statements of which include Varela, Thompson and Rosch (1991) and Thompson (2007). Enactivists in this tradition combine insights from phenomenology (in particular, Merleau-Ponty, 1945/2013) with work on the biological phenomenon of autopoiesis (e.g., Maturana and Varela, 1980). They claim that the interactive processes associated with cognition are continuous with those associated with life. Thompson observes that a cell, for example, produces its own material constituents through a metabolic interaction with the outside environment. Through this process of self-creation, the cell establishes the significance of certain environmental perturbations as either leaving the organism intact or resulting in its disintegration, and in this sense creates its own environmental niche. Living organisms arise out of self-organising biological processes like these in the way a tornado arises out of a dynamical process: they are, as Thompson puts it, invariant patterns that subsist in a changing physical substrate, implying they are not merely causally dependent on interactive processes, but constituted by them. Biological enactivists claim that self-organising biological processes of an appropriate kind are by nature minded, and that appealing to them helps close the explanatory gap between consciousness and the physical. To count as minded, an organism must be autonomous and answerable to norms. The appeal to biological self-organisation helps explain this. Organisms are autonomous by virtue of self-organising, and answerable to

¹ The label ‘embodied cognition’ is also sometimes adopted by quite different accounts, in which cognition is wholly brain-bound but involves neurally-realised representations of bodily states.

norms because their existence is predicated on responding to the environment in an appropriate way.

A recent entrant to the enactivist running is Hutto and Myin's (2012) 'radical enactivism'. In common with most statements of the sensorimotor theory, the authors exclude the biological apparatus distinctive of enactivism in its original form. Their distinctive contribution is to emphatically reject any theoretical role for representation or content, i.e. physical states bearing truth or correctness conditions about other physical states. In place of content, the authors claim that 'basic' cognition, meaning any cognition that is not shaped by natural language, consists of nothing more than interactive engagements between the brain, body and the external objects to which the cognition is intentionally-directed, which the authors propose can be explained by appeal to a nonrepresentational teleofunctionalist theoretical framework.

Noë gave the label 'enactive' to his (2004) version of the sensorimotor approach, not to declare a strict allegiance with biological enactivism, but to highlight its resemblance to it. Like biological enactivism, the sensorimotor theory downplays internal representation, and emphasises the role played by bodily interaction with the environment. In particular, it claims, like biological enactivism, that perception depends on the perceiver's mastering, in a bodily way, the character of its own bodily relation to the environment, in particular the way its orientation to objects in the environment would change given possible movements. Noë also expresses some sympathy toward the enactivist claim that consciousness requires life.² Thus the sensorimotor theory is typically regarded, with some justification, as a variety of enactivism. Hutto and Myin complain that the sensorimotor theory retains objectionable orthodox commitments on the ground that it retains a commitment to content and representation (although whether they are right about this is controversial, and Myin is himself, we should note, one of the original proponents of the sensorimotor theory). Despite criticising the sensorimotor theory for being insufficiently radical, the authors have, by including it in their taxonomy of enactive theories, helped crystallise the consensus that the sensorimotor theory deserves the 'enactivist' moniker.

² This claim does not, however, feature elsewhere in the sensorimotor theory, and is rejected by O'Regan, 2011.

1.2 - Structure of Thesis

In spite of its influential nature, the sensorimotor theory is formulated in differing and sometimes inconsistent ways throughout its canonical statements as well as the secondary literature, and in some cases its commitments have not been spelt out in sufficient detail at all. For instance, it is not clear whether the theory should, broadly speaking, be regarded as partly continuous with representationalist approaches in cognitive science, or allied with anti-representationalist embodied or enactive approaches such as those endorsed, for example, by Thompson (2008), Chemero (2009) and Hutto and Myin (2012). More narrowly, it is not clear what the sensorimotor skills the theory appeals to actually are, nor what kind of physical mechanism is implicated in them. Other examples of vaguely or inconsistently defined theses will become evident as the thesis progresses, and part of my aim is to identify aspects of the sensorimotor theory that are poorly defined and go some way toward clarifying them.

In chapter 2, I attempt a careful formulation of the negative claims that properly motivate the sensorimotor theory, an aspect of the account which is easily misunderstood and which I argue has been misunderstood, in one notable case, by Prinz (2006). Chapter 3 serves the dual purpose of laying out the sensorimotor theory's tenets in detail, and offering what I take to be the most promising and internally consistent ways of understanding them. Chapter 4 identifies some key limitations with the theory's notion of sensorimotor 'understanding', and offers a new proposal about what we should take sensorimotor understanding to be.

A later part of the thesis addresses some of the key criticisms that have been levelled against the theory, in particular by Clark (2008) and Block (O'Regan and Block, 2012). Chapter 5 addresses a worst case scenario in which it has been established that perception does in fact depend on nothing more than the activation of internal representations, which are not themselves intimately involved in the agent's sensorimotor engagements with the outside environment. I argue that in this event, there is still good reason to endorse the sensorimotor theory as an account of phenomenal qualities, and I offer some suggestions about how such an account could work. Chapter 6 offers a more robust defence of the sensorimotor theory in which I attempt to make it plausible that perceptual experience is constituted by the exercise of capacities for skillful bodily behaviour, and argue that representations should at most be

thought to play a heavily restricted theoretical role.

Chapter 7 moves beyond the core principles of the sensorimotor theory and begins to address new contributions the sensorimotor theory can make to long-discussed perceptual phenomena. Specifically, it addresses temporal extension. I argue that both the activity of perceiving and the contents of perceptual experience are temporally extended. Thus I suggest that object experience as well as event experience has a temporally-extended content, which temporally tracks the process of perceiving.

The aims of this thesis can be expressed as a unified goal. By clarifying some of the theory's founding principles, I aim to show how it can address some of the most prominent objections that have been levelled against it. By doing this, I hope to illustrate the novel contribution the sensorimotor approach can make to resolving some longstanding puzzles. Thus I hope to leave the sensorimotor theory in better shape than I found it.

2 - Making Conceptual Space

The sensorimotor approach makes a number of striking and relatively novel claims about perceptual consciousness and the physical processes that underlie it. Because of the more or less unorthodox nature of its positive claims, it is natural that the authors aim to motivate the approach by highlighting, with some emphasis, the faults and limitations of older approaches, and large early sections of O&N and Noë (2004) are devoted to this. The problem with this negative aspect of the sensorimotor theory is that many of the claims under target are not widely endorsed by contemporary theories, and in some cases merely reflect naive misconceptions or defunct historical ideas. In its standard presentation, the sensorimotor account is not pitched as a rejoinder just to mistakes of this sort, but also to the well-established orthodox view of vision in cognitive science, which holds that perception is realised by the construction and deployment, in the brain, of a detailed representation of the outside world.

The main part of this chapter will examine the manner in which the objections levelled in Noë and O'Regan's original statements of the sensorimotor approach cast doubt on the mature orthodox model. I will argue that they do not do enough by themselves to displace the orthodox view's core dogma that visual consciousness depends on internal representation, but they do, at least, indicate that orthodox theories must be careful to avoid certain mistakes. They may also intimate that the orthodox model is founded, implicitly, on a misplaced ambition, namely the desire to repair an obsolete 'inner movie' idea of vision that would better be replaced entirely. Because of the limited reach of the authors' negative claims, the sensorimotor account must be framed so that its appeal does not entirely depend on them.

2.1 - Kepler and Descartes: The Lessons from History

In a passage titled 'Fallacies Pictorial and Homuncular', Noë (2004) addresses historical accounts of vision, with a view to drawing morals for present day theory. Drawing on Lindberg's (1976) overview of pre-modern accounts, he calls attention, in particular, to Kepler's work on the optics of vision. Given what he knew about the eye's anatomy and the

principles that govern the refraction of light, Kepler was able to correctly identify that light meeting the lens is inverted (from top to bottom) and reversed (from left to right) before finally being projected onto the retina. This finding troubled Kepler. In addition to noting, quite appropriately, that the retinal image plays an important role in conscious vision, he thought that it did so by actually functioning as a picture. If the retinal image is inverted and reversed relative to the outside environment, it does not bear a pictorial resemblance to the environment as perceived. For this reason, Kepler ‘tortured’ himself attempting to identify a mechanism by which the light is re-inverted and re-reversed before reaching the retina. Finally abandoning this effort, he concluded that the disparity between visual experience and the projected image had to be corrected by a visual faculty occurring beyond the retina - involving, as he thought, spirits - and falling outside the scope of his own interest in vision, which was limited to optics.

Noë notes that Kepler’s misconceived search for a mechanism to re-invert and re-reverse the light was motivated by the more basic error, which Descartes later cautioned against, of supposing that the retinal image requires a further perceiver, a homunculus, to look at it. This is a faulty way of thinking, since it merely introduces the requirement for another homunculus to explain how the first one sees, and, following this pattern, an infinite regress of homunculi prevents explanatory closure. We should note that it is not obviously incoherent, in principle, to imagine that an image functions as a picture even in the absence of an observer, although claiming that a mere image constitutes a ‘picture’ requires some kind of explicit justification. Consider that a naturally occurring and function-less image such as a reflection of the sky in a lake does not merit being described as a picture in normal usage. Kepler, in any case, did think a homunculus was needed to view the picture, and for this reason was guilty of attempting to explain a psychological trait by introducing a homunculus that possesses the very same psychological trait - the most elementary version of what Kenny (1971) has called the ‘homunculus fallacy’.³ In spite of warning against the error, Descartes committed it himself by implying that the immaterial soul, serving as a homunculus, registers images presented to it at the pineal gland.

³ In section 2.1.3 I will address some different kinds of homuncularity, but this for now is the kind of homunculus under consideration.

2.1.1 - Pictures, Filling-In and the Bridge Locus

Noë (2004) claims that a confusion similar to the one that dogged Kepler persists into the present day. He gives the example of filling-in, a process by which the brain is alleged to compensate for information missing from the retinal image. Proponents of filling-in claim that the brain makes inferences about environmental features obscured, for example, by the retina's blind spot, and uses this inference-making mechanism to reconstruct the missing detail in an internal representation (e.g., Goodman, 1978).

On this topic, Noë, with reservations, endorses Dennett's (1991; 1992) position. Dennett is known for criticising the view that the brain's visual processing is geared toward the production of a Cartesian theatre, an end-product which occurs in a spatially circumscribed area of the brain and constitutes the necessary and sufficient substrate of visual consciousness. Dennett claims that this view, 'Cartesian materialism', is motivated by a tacit commitment to the existence of a homunculus for whose benefit the visual processing takes place. This is a problematic idea for the reason I have just noted, although Dennett is not concerned with homuncular regress so much as with implausible theoretical views that follow from positing a homunculus, the full details of which need not be rehearsed here. He observes that Cartesian materialists are forced to presume that the brain fills-in data missing at the blind spot because they imagine that the notional homunculus would otherwise notice that the information is missing. Because he rejects this position, Dennett claims instead that the brain simply ignores the lack of information at the blind spot. Noë (2004) enthusiastically endorses Dennett's rejection of Cartesian materialism, and in diluted form the scepticism about filling-in. He makes the important proviso that filling-in is only a misguided posit if the need for it is assumed *a priori* as a result of prior commitment to the existence of a homunculus. This does not rule out the possibility that the existence of filling-in can be established by appeal to relevant empirical work.

Visual stability is a further domain in which Noë suggests contemporary theories are sometimes guilty of the homunculus fallacy. It is sometimes considered noteworthy that despite continual saccadic eye movements and attendant changes to retinal input, the world is presented in visual experience as being stable. There is a temptation to assume that the brain

must deploy a special mechanism to compensate for this, as if there were a homunculus who would otherwise confuse the resultant changes to the positions of objects in the retinal image for changes in the spatial position of the objects themselves. Noë argues that a compensatory mechanism is only a conceptual necessity if you think that perception consists of a pictorial representation derived from the retinal image and offered up for the benefit of a homunculus. If you do not think this, you need not imagine that visual experience is in the first place prone, without a special compensatory mechanism, to mistaking eye movement-related changes in retinal input for changes in the outside world. Thus, an *a priori* commitment to the existence of the compensatory mechanisms only arises if you commit the homunculus fallacy. O&N and O'Regan (1992) make a similar point concerning saccadic suppression. Saccadic eye movements create a motion blur effect on the retinal image, and it is sometimes thought that the brain inhibits informational input from the retina during saccades in order to prevent perceivers from experiencing a blurred visual scene (Matin, 1974). O&N suggest that this is likely to be motivated by a commitment to a homunculus, since the smear effect need not be compensated for if the retinal image does not function as a picture.

These examples aside, Noë is imprecise about the intended scope of the historical lesson. He claims that: “No contemporary theorist believes that we see by seeing internal pictures” (2004, p. 46), but finishes the passage by stating: “The purpose of this section has been to reveal the extent to which our thinking about perception, like that of Leonardo and Kepler, is tied to a problematic conception of the need for pictures in the head” (p. 48). What Noë means when he says that theories are *tied to* this problematic conception is ambiguous, given that there are many ways a theory could be ‘tied’ to some conception which need not involve explicitly endorsing it.

Prinz (2006), responding specifically to *Action in Perception*, takes it that Noë’s intention is to draw a rough and ready parallel between Kepler and a contemporary version of the ‘pictorial’ view, in which neural states yield visual awareness with a particular content by virtue of their pictorial qualities. Prinz observes that: “We do not have color filled pixels in the head, much less an inner observer to discern all the parts. But this view is not seriously defended by anyone” (2006, p. 12). Appealing to Kosslyn (1994) as more representative example of a contemporary pictorial theory, he states that Kosslyn takes care not to require an

infinite regress of homunculi, and observes that pictorial views in general are not pre-committed to the existence of filling-in at the blind spot, since they could instead endorse Dennett's suggestion that the subject simply ignores the absences of information.

Prinz continues by citing a study by Ramachandran and Gregory (1991) in which subjects were presented with a television screen showing at first a 'snow' pattern and subsequently switched to show a solid colour instead. Subjects reported undergoing a briefly continued awareness of the snow pattern in an area of the visual field corresponding to the blind spot, even after the stimulus had been changed. The study suggests that information missing at the blind spot is not, in fact, simply ignored in visual processing, as if it were there would be no way to explain the continued percept of the snow in this region of the visual field. Prinz claims that the evidence that filling-in does take place counts as an embarrassment for the sensorimotor theory.

Noë is admittedly vague about the lesson he intends us to take from Kepler and Descartes, and a sceptic would be well-justified in observing that Noë does not comprehensively address the numerous resources that orthodox, including pictorial, accounts appear to have to avoid the homunculus fallacy. However, Prinz's rebuttal falls wide of the mark. Noë neither claims that pictorial accounts necessarily commit the homunculus fallacy and therefore fail, nor denies the existence of neural filling-in. A better account of his position can be found in an earlier paper by Pessoa, Thompson and Noë (1998). The paper rejects Cartesian materialism and also a doctrine the authors call 'analytic isomorphism', the *a priori* thesis that the substrate of perceptual consciousness must be a circumscribed area of the brain - a 'bridge locus' (Teller and Pugh, 1983) - which bears a functional (including topological) or topographical (i.e. map-like) isomorphic correspondence to the features represented by the subject's experience. Cartesian materialism and analytic isomorphism resemble one another in that they are both committed to the existence of a bridge locus, although isomorphic correspondence between the neural and the perceptual is in principle not essential to Cartesian materialism, and proponents of analytic isomorphism could deny that their position is motivated by any tacit or explicit commitment to the existence of a homunculus.

Pessoa et al. claim that if you endorse analytic isomorphism, you have a different reason to assume *a priori* the existence of filling-in. To maintain the isomorphism between the

bridge locus and the content of perceptual experience, there must be a neural state which is isomorphic to features that the experience represents as occurring in the location obscured by the blind spot, given that the subject does not experience the blind spot as an absence. The authors take it that once we reject Cartesian materialism and analytic isomorphism, there is no reason to aprioristically assume the existence of filling-in, although they actively argue for the claim that filling-in does take place, appealing to empirical data which I will not rehearse given that there is no disagreement on this point between the authors and Prinz.

A similar point holds with regard to pictures in the head. The pictorial view states that image-like physical states help generate visual awareness by serving as pictures, and this implies they are topographically or topologically isomorphic to the features represented in the experience, regardless of whether or not there is supposed to be a homunculus to look at them. Thus, pictorial theories could fall under the doctrine of analytic isomorphism, although they need not. Pessoa et al. are explicit, however, that they are open, subject to empirical evidence, to the claim that neural states help enable visual consciousness by virtue of bearing a pictorial resemblance to the features represented in visual consciousness. The authors reject the aprioristic commitment to isomorphism, along with the suggestion that where pictorial representations do play a role, this takes place in a spatially localised neural region which serves as a bridge locus. Noë (2004) says nothing to contradict any of the claims I have just mentioned from the (1998) co-authored paper.

Notice that for each of the phenomena Noë (2004) and Pessoa et al. (1998) target - pictorial representations, filling-in, compensatory mechanisms for visual stability, and neural-perceptual isomorphisms - the claim is not that these phenomena do not exist. The authors are careful to claim that the phenomena may exist, and in some cases probably do. A better moral is that confusion arises when we make theoretical commitments to Cartesian materialism or analytic isomorphism necessitating the existence of something that should be an empirical posit, a diagnosis that at least holds true when we think of Kepler's commitment to pictorial representations and a homunculus. In one sense, then, the negative claims motivating the sensorimotor theory are more modest than Prinz suggests. Noë should not be taken to be attempting to actively disprove pictorial theories or other representational theories that, like Kosslyn's, take care not to commit the homunculus fallacy.

Prinz's other objection is that Noë is targeting a straw man, given that no one seriously defends the claim that we have colour-filled pixels in the head or a perceiver to discern them. Strictly, this may be true. Pessoa et al. observe, however, that in the commentaries to their *BBS* paper, Lehar (1998) explicitly claims that spatial perception logically requires the existence of a three-dimensional neural model. In his own widely-cited *BBS* paper, Lehar (2003) claims that topographical or topological isomorphism is not a logical requirement for perception, but that functional isomorphism is. Using the (2003) paper as a guide, Lehar's requirement can be reconstructed in the following way. Imagine that the brain contains a 3D model in the shape of a cube and that this correlates with the conscious percept of an identically-shaped cube. The model is thus topographically isomorphic to the perceived cube. Now suppose that instead of being perfectly cube-shaped, the model has the shape of a stretched-out cube, with some planes that are oblong rather than square. The model continues to be topologically isomorphic to the perceived cube, because the model's planes continue to stand in appropriate spatial relations to one another. Informally, this is because a cube could be distorted into a stretched-out cube shape without ripping or tearing.

Lehar claims that topological or topographical isomorphisms are likely to play a role in perception, but are not required as a matter of logical necessity. Now, imagine that instead of merely being stretched-out, the cube-shaped neural model is scrambled at random, such that it no longer looks anything like a cube, even a distorted one. This means it is no longer topologically or topographically isomorphic to the perceived cube. Think of the scrambling, however, as a warping of the geometry of the space the shape inhabits, rather than merely a warping of the shape itself. Suppose that the perceived cube is transformed in some way, for example rotated. If the model is transformed in the way that a topographical model would have been, allowing for the distorted geometry of the space it inhabits, a functional isomorphism obtains between the model and the perceived cube.⁴

Lehar insists that functional isomorphism, at least, is logically necessary for perception. This is a strikingly robust claim, and almost sufficient for the theory to constitute an instance of analytic isomorphism, although the appeal to functional isomorphism avoids the logical

⁴ Pessoa et al. (1998) use the term 'topological isomorphism' to refer to topological *or* functional isomorphism in Lehar's usage. I am following Lehar's usage.

requirement for a localised bridge locus. Lehar suggests, however, that there is reason to think, prior to empirical confirmation, that topographical and topological isomorphisms are particularly likely to obtain, since they are more efficient than other functional isomorphisms to implement. This is a demanding theoretical commitment, and it suggests that Lehar's view is biased toward the existence of a bridge locus, even if it does not require it, since these allegedly more efficient forms of isomorphism cannot be distributed across the brain in the way functional isomorphisms could be. This shows that analytic isomorphism is not far removed from explicit claims found in recent and influential work.

Lehar's view is not a mainstream position, however: A moral more widely applicable to representationalist cognitive science is that we should be deliberate about rejecting any aprioristic commitment to posits that are only needed if one does endorse analytic isomorphism or Cartesian materialism. Notice that doing so opens up conceptual space for a range of posits that, unlike pictorial representations and filling-in, are actively incompatible with those commitments; and some of the alternative views thereby made available may prove persuasive on positive rather than negative grounds. After rejecting Cartesian materialism, Dennett claims that consciousness depends instead on representational processes that are widely distributed across the brain. The sensorimotor theory does something similar, but makes a bolder positive proposal by claiming that cognitive processes sometimes depend in a non-trivial and frequently unacknowledged way on the subject's embeddedness within the outside environment and, on some formulations, that the substrate of perceptual experience actually includes activity in the extra-neural body and outside environment.

I will elaborate on the sensorimotor theory's positive claims in chapter 3, but will here briefly highlight the sensorimotor theory's alternative proposals with regard to filling-in. Pessoa et al. observe that rejecting analytic isomorphism and Cartesian materialism allows for a pluralistic explanation of perceptual completion, the visual phenomenology sometimes accounted for by filling-in. They propose that in some cases this involves pictorial representation, in other cases symbolic representation, and in others requires the absence of information to be ignored in the way Dennett suggests. Their account also holds open the possibility that perceptual completion might in some instances be explained by appeal to O'Regan's (1992) idea of the world as an 'outside memory'. O'Regan's idea is that perceptual

experience depends on a temporally-extended, embodied process in which perceivers gain access to missing detail by attending to different parts of the visual field, making use, in particular, of eye movements to access missing detail as needed directly from the outside environment without needing to 're'-present that information internally (see section 3.1).

2.1.2 - Neural Representation

I now want to offer another lesson we can learn from the historical discussion featured in Noë (2004), in a similar spirit to arguments made by Noë (2004) and Pessoa et al. (1998) but more unorthodox in its implications. Just as it would be wrong to assume, without good evidence, that filling-in takes place, it would also be a prejudice to commit oneself in advance to the claim that the neural processes involved in perception - including any inner processes that compensate for information missing from the retina - must involve representation. Some theorists (e.g., Wheeler, 2005) claim that internal representation requires homuncularity, although not necessarily of a problematic kind. I endorse this claim in section 2.1.3. The important point for now, however, is that to assume out of prejudice that the neural states involved in perception involve representation is to betray a commitment to the kind of homunculus which Descartes warned against. Descartes observed that retinal images play a causal role in perception, but that you are likely to be in the grip of the homunculus fallacy if you conceive of them as 'pictures'. By the same token, the states and processes involved in perception, including those involved in filling-in, might play a causal role without functioning as representations. It would be a mistake to pre-theoretically commit oneself to the existence of representation when the neural activity that helps enable perceptual experience might be better conceptualised as a nonrepresentational causal process.

Deciding whether the neural states involved in perception, including filling-in, are representations is in part an empirical enterprise. Indeed, it is possible that appeals to internal representation could be justified on entirely empirical grounds. Suppose that there were a unified structure in the brain that looks like an expansive and finely-detailed three-dimensional model of the visual scene, and we could use this to accurately predict the character of the subject's reported perceptual state. Few theorists realistically entertain this,

and even Lehar (2003) falls short of insisting on it, but it remains a conceptual possibility. If this were the case, it would be reasonable to suppose that the neural structure is a 'model' or 'representation', even in the absence of a philosophically robust explanation of how the structure comes not just to resemble the world but, strictly speaking, to 'represent' it.

A similar possibility holds true even if the purported neural representations are not supposed to involve neural-perceptual isomorphisms. For instance, it might turn out that we can accurately predict a subject's perceptual state on the basis of the subject's present neural state in combination with a theory of vision that makes reference to representation purely because appealing to representation is the most useful way to explain the neural activity. The idea that a neural state can be legitimately described as a representation purely because it is expedient to do so is widely endorsed, even by non-representational theorists like Chemero (2009), although it is rejected by Hutto and Myin (2012) - more on which in section 6.4.

If we agree that there is filling-in, this does not licence the claim that there is necessarily representation. There may be processes that complete visual information missing from the retina and in so doing affect the character of visual consciousness. However, should we abandon the view that there is a detailed inner world model (a kind of representation) to begin with, we abandon another important motivation for insisting that these processes necessarily constitute representation, even of a non-pictorial kind. The suggestion that there are processes that resemble filling-in without representation should be no more controversial than Descartes' assertion that the retinal image plays an important causal role without functioning as a picture. If we construe the import of the Kepler discussion for the sensorimotor theory this way, we have another reason for thinking that Prinz's (2006) contention that empirical evidence for filling-in should embarrass the sensorimotor account is wide of its mark.

It is an empirical bet, however, that neural states best understood as representations do exist. If neural representations do exist, it is also an empirical bet that appealing to their content or functional role is sufficient by itself to explain the target phenomena. For example, we might find that some neural process plays a functional role equivalent to filling-in, but also find that the process, taken together with other neural processes explicable in terms of representation, are insufficient to show which features the perceiver is intentionally-directed

toward at the agent level, or to account for the attendant phenomenal character (if the intentional directedness and phenomenal character of an experience could come apart, a possibility that may not be coherent). Pessoa et al. (1998) and Noë and Thompson (2004) deny that subpersonal content can be used to predict personal level perceptual states in this way, although they do not deny that subpersonal representation exists.⁵ If it proved impossible, however, to explain or predict the intentional or phenomenal character of the subject's experience without pointing to non-neural features, including, for example, variations in the agent's behavioural capacities, this would deprive representationalists of a significant reason for thinking that perception is realised by internal representations at all.

The point of this passage has not been to claim that perception does not involve representation, but, as before, to make conceptual space for the possibility that internal representation might play a reduced or eliminated explanatory role in a final account of perception and phenomenal consciousness. Many who are open to representationalist

⁵ Personal-level mental states are those it is appropriate to attribute to a whole person, and are generally thought to necessarily be conscious or at least consciously accessible. Subpersonal states are those which it is appropriate to attribute to a person's parts. For example, a neural representation may bear content about some state of affairs, even though the conscious subject herself does not entertain mental states with that content. Some states may be subpersonal and personal, just in case they can be intelligibly attributed to a subpersonal subsystem but are also consciously accessible. Noë (2004) expresses doubt about the value of the distinction on the ground that certain activities, such as eye movements, sometimes come under rational control by the person, and sometimes do not, and that both cases are likely to be similarly important to perception. What he has in mind is McDowell's (1994b) use of the distinction, which Drayson (2014) implies is an attempt by the Pittsburgh school of philosophers to hijack the distinction for their own idiosyncratic purposes. Personal level states on this perspective are those occurring in the space of *reasons*, meaning they are properly considered to be mental states, while subpersonal states are those occurring in the space of *causes*, the latter category denoting states that are not genuinely mental states at all but non-minded physical states. Note that on the non-Pittsburgh understanding, the fact that subpersonal states are sometimes assimilated to the personal level is to be expected, meaning that Noë's point does not necessarily compromise the distinction in its original meaning. In any case, Noë's sole-author account is marked by notable shift from talking about brains and artifacts, as O&N did, to talking about whole persons. Moreover, he endorses a distinction between the 'constitutive' and the merely 'enabling', which in McDowell's paper tracks the distinction between the personal and subpersonal, and which may have a significance analogous to the personal/subpersonal distinction in Noë's account. I explain this latter distinction and examine its significance later in the thesis.

accounts of cognition endorse the suggestion that the legitimacy of internal representation as a theoretical posit cannot just be assumed, but must be justified relative to the particular theory it is deployed in (see Ramsey, 2007, for a detailed defence of this point). However, others who are critical of non-representational accounts are prone to claiming that cognition, including perception, *just is* a representation-involving process.⁶

2.1.3 - Varieties of Homuncularity

I have, up until now, been discussing the most elementary version of the homunculus fallacy, identified by Descartes when he rejects the suggestion that retinal images function as pictures, warning: “We must not think that it is by means of this resemblance that the picture makes us aware of the objects - as though we had another pair of eyes to see it, inside the brain” (1637/1985a, p. 167). Ironically, contemporary theories that make this kind of mistake are described by Dennett (1991), pejoratively, as examples of Cartesian materialism, on the ground that Descartes went on to make the same mistake by claiming that the immaterial soul is presented with images at the pineal gland.

This elementary version of the fallacy is considered by some (although not Dennett) to be just one instance of a wider fallacy, allegedly more pervasive than mere Cartesian materialism in present day cognitive science. The error, sometimes known as the mereological fallacy, can be defined, broadly, as the ascription to parts of things predicates that should properly only be ascribed to wholes. Narrowly, it is the ascription of psychological predicates, for example ‘seeing’ or ‘believing’, to parts of persons - in particular, brains - where they should only be ascribed to persons, i.e. the complete living organism (Kenny, 1971; Bennett and Hacker, 2003). In the basic version of the fallacy, applying a psychological predicate to the brain is uncontroversially a mistake. If we try to explain the faculty of vision by positing the existence of a homunculus (a part of the person) that possesses an identical faculty, we find an infinite regress ensues of the kind discussed earlier. It is telling that the fault, in this instance,

⁶ For proof that many theorists take cognition to involve representation *by definition*, see Hutto, Kirchhoff and Myin (2014), who offer textual evidence from Khalidi (2007) and O'Brien and Opie (2009) to this effect.

lies in the circularity of the proposed explanation.

A.O. Rorty (1971), replying to Kenny (1971), notes that circularity of this sort only occurs if you try to explain a person's psychological state by ascribing to their brain or parts thereof *exactly the same* psychological state. The threat of circularity does not indicate the existence of a more general problem with ascribing person-like psychological predicates such as 'believing' to brains, since circularity is avoided so long as the explanandum and explanans feature different psychological states. Dennett (1978) is mindful that a certain kind of circularity may persist, however, even if the intentional predicates attributed to the part and the whole are different, or attributed to different degrees. To account for individual faculties like thought or vision, we need to explain, more generally, how intentionally-directed psychological states can be realised by a physical system at all. To account for a psychological attribute possessed by the person by reference to psychological attributes possessed by the brain is to presuppose that something bears intentionality, even if the states or processes featured at each end of the explanation are different. Dennett therefore proposes, like Rorty, that a person with complex psychological attributes might subdivide or 'decompose' into parts - i.e. homunculi or 'subpersons' - possessing simpler psychological attributes. Those homunculi might, in turn, decompose into still simpler homunculi. To show how psychological attributes might be realised in the first place, Dennett introduces the crucial suggestion that the bottom layer of homunculi are so simple, they decompose or 'discharge' into physical processes that can easily be described using non-intentionalistic vocabulary.

For some, however, circularity of the sort I have been discussing is merely a symptom of a more basic error, a variety of category mistake. Bennett and Hacker (2003) and Kenny (1971) both quote, approvingly, Wittgenstein's assertion that: "Only of a human and what resembles (behaves like) a living human being can one say: it has sensations; it sees; is blind; hears; is deaf; is conscious or unconscious" (1953/2009, para 281). The moral, for Kenny, is that since it is not *obvious* that brains resemble humans in the relevant way, we should proceed with caution, rather than unthinkingly ascribing the properties actually possessed by the whole person to parts thereof. Bennett and Hacker (2003) draw a more demanding conclusion. They argue that it is evident from everyday linguistic convention that psychological or otherwise intentionally-loaded vocabulary such as 'representing', 'mapping',

‘understanding’ and ‘believing’ can only be meaningfully used in relation to persons rather than brains. On this basis, and irrespective of any particular empirical facts about brains, Bennett and Hacker claim that neuroscientists, psychologists and cognitive scientists are guilty of conceptual incoherence when they apply these predicates to the operation of brains.

While homuncular decomposition offers, in principle, a means to avoid circularity, it does not, in itself, address the more fundamental charge that brains and parts of brains simply do not, in the relevant way, resemble or behave like humans. Dennett (2007), replying to Bennett and Hacker, remarks that we cannot use ordinary linguistic practice alone as a guide to the applicability of intentionalistic predicates to subpersons, since the claim that we can use these predicates turns on the empirical thesis that subpersons do, indeed, behave like human beings. Dennett is surely right that the admissibility of homuncular decomposition depends not just on aprioristic conceptual analysis but also on empirical facts about brains. There is a pressing uncertainty, however, about the sort of properties or behaviours that would need to be identified through empirical enquiry to show that an object is sufficiently person-like to licence the ascription of psychological predicates.

Dennett (1987) offers a specific proposal that could answer this. According to the ‘intentional stance’, we can legitimately attribute psychological predicates like beliefs and desires to an object, including a person, creature or artifact, just in case ascribing them is a useful way to explain and predict the object’s behaviour. The account is explicitly intended to be liberal enough to allow us to ascribe intentionality not just to people but non-living artifacts like thermostats and chess-playing computers.⁷ Dennett asserts not only that the intentional stance is scientifically respectable, but also that it reflects everyday folk

⁷ Dennett (1987) does not actually claim that we should apply the intentional stance to brains. He suggests that at this level we should apply the design stance, meaning we conceive of them as being designed to achieve particular functions. However, the homuncular decomposition strategy requires us to view brains and parts of brains as person-like, meaning the design stance does not appear to do the job. The intentional stance actually undercuts one of the roles homuncular decomposition might play, namely to resolve the fundamental question - a sort of hard problem - about how mere physical states could be intentionally-directed. The intentional stance solves this problem without requiring homuncular decomposition. However, endorsers of the intentional stance may still usefully make use of the homuncular decomposition strategy to show how intentionalistic descriptions can finally be cashed out using non-intentional, physicalistic vocabulary.

psychological practice, noting that in everyday conversation we frequently use belief and desire talk, for example, to make sense of behaviour by machines. If this is correct, then it is not implausible that brains, and parts of brains, can be ascribed intentionality in this way, too. In this event, the mereological fallacy is no threat to homuncular decomposition, either by the standard of everyday linguistic practice *or* the standard demanded by our theoretical account of intentionality.

Bennett and Hacker (2003) are on highly contentious ground when they suggest that everyday usage gives us a determinate guide to the correct ascription of intentional predicates, since one might think that we do not necessarily know precisely which features make our everyday ascriptions of psychological predicates true or false. It might be that in folk psychological discourse, we are employing the intentional stance. In this event, Bennett and Hacker are hoist with their own petard, for the reasons I have just suggested. Alternatively, it may be that everyday use of intentionalistic language picks out, unbeknown to us, properties unique to people or other complete living organisms, meaning the ascription to intentionalistic predicates to brains or artifacts are homonymic, metaphorical⁸ or just false. Bennett and Hacker's argument is enough to show that endorsers of homuncular decomposition have a positive need to endorse an account of intentionality that makes the approach tenable.

Depending on the approach to intentionality you endorse, the empirical criteria required to establish the existence of 'homunculi' may be more or less demanding. Adopting the intentional stance, homuncular decomposition seems relatively unproblematic. Alternatively, if you endorsed, for example, the autopoietic enactivist thesis that only a certain kind of autonomous biological agent can manifest intentionality (e.g., Thompson, 2007), homuncularity would appear much less tenable. Similarly, those representationalists that suppose intentionality is best explained by robustly naturalised subpersonal content (e.g., Millikan, 1984; see section 6.4) must maintain that the homunculi endorsed as 'top down'

⁸ McDowell (1994b), for instance, argues that content-talk at the subpersonal level is metaphorical, but nonetheless legitimate. In chapter 7, I argue that metaphorical representations and content is legitimate in principle, but may not be useful or appropriate in practice.

theoretical posits by Dennett's approach find counterparts in evolved modular architecture concretely implemented by the brain, a point made by Wheeler (2005). In either of the latter two cases, the burden of empirical evidence required to justify ascribing psychological predicates to parts of the brain will be strong.

I will not attempt to rule, now, on the correct account of intentionality. I will conclude this section by noting, however, that while Bennett and Hacker's automatic dismissal of intentionalistic vocabulary at the subpersonal level is misplaced, endorsers of this vocabulary must offer an account of intentionality that is amenable to the existence of homuncular decomposition, naturalistically plausible, and, we might think, motivated by considerations other than the desire to offer a *post hoc* justification for their use of intentionalistic language. This is a heavy burden, and it may prove that a more promising approach is to reformulate the vocabulary in question instead.

2.2 - Conservative and Radical Strands

As we have seen, the negative claims that properly motivate the sensorimotor theory are geared toward unseating a number of prejudices rather than actively ruling out the suggestion that perception depends on nothing more than the construction and deployment of internal representations, even representations that are richly detailed, pictorial or realised by a circumscribed area of the brain. However, the arguments are sufficient to create a legitimate suspicion that conceptions of vision endorsing these representationalist posits are motivated, at least in part, by a residual commitment to faulty historical assumptions, even if they have resources to avoid the homunculus fallacy.

The fact that the negative claims are relatively modest means there is leeway in how radical we take the theory's positive account to be. I will continue by outlining two competing versions of the sensorimotor approach, one *conservative* and one *radical*. Although they are potentially in significant disagreement with one another, the competing accounts reflect, I will argue, an uncertain attitude to representation present in the joint and sole-authored work of both Noë and O'Regan. The conservative account embraces a moderate construal of the authors' negative claims, and consequently endorses the orthodox thesis that perceptual

consciousness is realised solely by the brain (the claim known as *vehicle internalism*) through the construction and deployment of internal representations (*representationalism*); by virtue of this concession, the account can be thought of, essentially, as an attempt to progress the orthodox model in a way that inoculates it against Cartesian materialism, rather than an attempt to radically break from orthodoxy. This conflicts with the sensorimotor theory's frequently radical tenor. However, the conservative approach continues to make a number of interesting claims, including the view that perceiving depends on active, goal-directed interrogation of the outside environment.

Many of the theory's claims place a distinctive emphasis on embodiment: Although the conservative sensorimotor approach denies that perceptual consciousness supervenes on extra-neural activity, it recognises that bodily exploration plays a non-trivial and generally unacknowledged causal role in the construction of the neural states that do form the supervenience base. On this basis, it asserts that bodily exploration must be mentioned in any adequate account of vision. Further, the conservative approach endorses the particularly distinctive thesis that phenomenal qualities, rendered mysterious on rival accounts, are best explained by appeal to laws of sensorimotor contingency, as I will explain. This, again, means our account of perception must refer to characteristics of the perceiver's body, specifically its morphology and the nature of the body-environment relation that attends. By virtue of the body and environment's essential explanatory role, the sensorimotor approach could, on this construction, be described as a form of 'embedded' and conservatively 'embodied' cognition. Given the limited scope of the theory's negative claims, the relatively modest ambition of the conservative approach is an advantage, since it only needs to upset that claim that perceptual awareness depends on richly detailed, pictorial representations, and not the claim that perceptual awareness supervenes on internal representations in general.

An alternative construction situates the sensorimotor theory near the radical end of a continuum of 'embodied' and 'enactive' approaches to cognition. Accounts in this region frequently deny that cognitive processing is realised solely by neural states. To this effect, they sometimes suggest that the representational vehicles which realise cognition frequently incorporate bodily or environmental activity (Clark and Chalmers, 1998; Rowlands, 2006). Alternatively, they deny that there is representational content or anything else that would

motivate a principled insistence that the boundaries between brain and body, or body and world, are of theoretical relevance in determining the bounds of cognitive processing (Hutto and Myin, 2012). Accounts at the radical end of the spectrum, especially those that adopt the ‘enactivist’ label, are also known for suggesting that there is no internal representation whatsoever (e.g., Varela, Thompson and Rosch, 1991; Chemero, 2009; Hutto and Myin, 2012).

2.2.1 - The Tenets of Sensorimotor Theory: A Brief Outline

In this section I will proceed by outlining the main tenets of sensorimotor theory, categorising them into claims endorsed by both conservative and radical variants of the approach, and claims endorsed solely by one variant or another. I do not mean to claim that any version of the sensorimotor account must endorse every one of the claims falling under either the conservative or the radical strand. However, identifying these distinct variants offers a useful way to examine differences in the explicit claims made by the canonical accounts of the sensorimotor approach, different ways of reading those accounts, and disagreements in other literature on sensorimotor theory, a number of which will be investigated as the thesis progresses. The first four claims are endorsed by both variants of the theory and the claims endorsed only by one or the other follow.

(1) *The World as an Outside Memory*. This claim draws on work by O’Regan (1992) and ‘active vision’ approaches in computer vision (e.g., Ballard, Hayhoe, Pook and Rao, 1997). Endorsing Brooks’s (1991) suggestion, in robotics, that the world serves as its ‘own best model’, the outside memory thesis (O&N) denies that human visual systems make use of a detailed inner world model, claiming instead that our visual systems access detail directly from the environment as required in the course of bodily exploration. This claim is given weight by empirical work on change blindness and inattention blindness, where it has been found that human subjects, with their attention drawn elsewhere, frequently fail apparently to notice occurrences that take place seemingly in full view. The outside memory thesis entails that visual processing is, at least, causally scaffolded by bodily exploration in an interesting and frequently unacknowledged way. It does not necessarily imply the metaphysical thesis that

perceptual consciousness actually supervenes on extra-neural activity, although the outside memory thesis may be a key aspect of an account that does make this claim.

(2) *The Sensorimotor Contingency Account of Phenomenal Qualities.* Phenomenal qualities (or ‘qualia’), for example the feel of a sponge, the look of red or the sound of a bell (O’Regan, 2011), are not irreducible. They can be accounted for, exhaustively, by reference to characteristics of the conscious perceiver’s bodily relation with the outside environment - specifically, sensorimotor contingencies (SMCs). In O&N, SMCs characterise the law-like patterns of dependence that hold between motor output signals and informational input from the sense organs. That paper does not specify explicitly, but they might be counterfactual dependences, expressible by statements of the form ‘if I moved like this, sense input *would* be like this’. More restrictively, they might be patterns evident in the bodily interactions that the subject is right now engaged in, for example the way a sponge they are presently squeezing between their fingers yields easily to touch. Burhmann, Di Paolo and Barandiaran (2013) note that the latter kind of SMCs depend not just on peripheral sense input and motor output states, but on the internal states of the agent, since they depend on how the subject’s motor responses vary in line with different sense inputs, and this depends on the brain. The account of colour offered by many statements of the sensorimotor theory (O’Regan, 2011; Degenaar and Myin, 2013) also appeals to patterns of counterfactual dependence outside the skin, since an object’s experienced colour depends on the ways sense input changes with movement and in line with changes to ambient lighting conditions. Finally, in Noë (2004), SMCs also describe perspectival properties, the ways objects consciously appear when considered from a subject’s own perspective (for example, the elliptical appearance of a plate when viewed at a non-perpendicular angle) and the ways these change in line with movement.

(3) *The Sensorimotor Mastery Account of Perception.* Perceiving takes place when the subject possesses or exercises ‘knowledge’, ‘mastery’ or ‘attunement’ to SMCs. Sensorimotor ‘mastery’, as I use it throughout, is a generic term intended to incorporate a number of loosely related notions found throughout the literature on sensorimotor theory, and intended to be liberal enough to incorporate a ‘radical enactivist’ version of sensorimotor theory. The proponents of radical enactivism, Hutto and Myin (2012), dislike the terms sensorimotor ‘understanding’ and ‘knowledge’ on the ground that these notions appear to

require representational content. Myin and Degenaar (2014) endorse, by way of developing a positive radical enactive gloss, a notion of ‘attunement’ instead, which I include under the umbrella of sensorimotor mastery (note, however, that their concept of sensorimotor ‘attunement’ is different to the one I develop in chapter 3). O&N and O’Regan (2011) refer to a variety of mastery which appears, like attunement, to denote a capacity for skillful bodily behaviour that does not necessarily call upon internal representation. The same accounts suggest, however, that *conscious* perception requires sensorimotor ‘knowledge’, which might be taken to denote a distinct kind of faculty.

(4) *Perceptual Presence as Access*. This idea, endorsed by Noë, is a natural counterpart to the outside memory claim. It suggests that the environment, in all its detail and expansiveness, has a felt phenomenological ‘presence’ to the conscious perceiver which outstrips the information processed, at any one time, by the perceiver’s visual system. This is explained by the perceiver’s possession of sensorimotor mastery, which in this case amounts to knowledge of (or perhaps attunement to) the movements s/he must make to find out what is present and, perhaps, an implicit expectation about what the sensory results of those movements will be. The ‘presence-as-access’ (or ‘virtual presence’) the environment enjoys in the perceiver’s conscious awareness is analogous to the presence to your computer of the web edition of the New York Times. When you read an online newspaper, you generally download just one article at a time. However, since you can access any of the day’s articles the moment you require them, they are, practically speaking, just as ‘present’ as they would be if you downloaded the day’s edition all at once (Noë, 2004).

The following claims are endorsed only by the radical sensorimotor approach:

(5) *Virtuality All-The-Way-In*. The presence-as-access thesis, above, could be straightforwardly endorsed by most sensorimotor theorists as a natural phenomenological counterpart to the outside memory thesis. Noë makes a more challenging proposal, however, when he says that in a disanalogy with the computer case, perceptual presence in humans is virtual “*all the way in*” (2004, p. 134). This means that there is no distinction to be found in the phenomenal character of perceptual consciousness between detail that is now present

‘simpliciter’ and detail that is present by virtue of being accessible in the manner described by the presence-as-access and outside memory theses. Noë claims that this suggestion finds confirmation in the introspectable phenomenal character of one’s own visual consciousness, alleging that you can never take in, all at once, any complete or atomic visual feature, for example, “*that shade of red*” (p. 134). If the virtuality *all-the-way-in* thesis is considered primarily as an extension of the outside memory thesis, it might seem to amount to a denial that perceptual consciousness exists or has any phenomenal character. As I will later argue, this impression can be avoided if the claim is considered, instead, primarily as a counterpart to the next thesis.

(6) *Perception as Dynamical Entanglement*. This claim is endorsed by Noë (2004) and reflects, in particular, important work by Hurley (1998). It involves rejecting the suggestion, endorsed by many classical accounts of cognition (see, e.g., Fodor, 1983), that there are functionally distinct and informationally encapsulated ‘input’ and ‘output’ systems where the activity of input systems has no tight causal link with the activity of output systems, or vice versa. Endorsers of the dynamical entanglement thesis claim, instead, that there is an intimate causal link, running in both directions, between neural activity in the sense input and motor output areas of the brain. Further, the thesis claims that there is no theoretically relevant distinction between activity in those areas and outside, in the extra-neural body and environment. As a result, perception must be explained by reference to patterns of dynamic activity occurring across the entire loop of interaction occurring between brain, extra-neural body and world.

(7) *The Extensive Conscious Mind*. The dynamical entanglement thesis is one of the key claims motivating the closely allied thesis, endorsed by Noë (2004) and Hurley (1998), that the supervenience base of perceptual consciousness literally extends beyond the boundaries of the brain, incorporating aspects of the extra-neural body and outside environment. The argument, briefly put, is that since perceptual consciousness depends on a dynamical system incorporating these extra-neural features, it may not be possible or appropriate to distinguish between the contribution made by these features and the brain when determining the substrate. This thesis is sometimes described as ‘vehicle externalism’ about consciousness or the ‘extended conscious mind’, although for reasons that will be

explained in section 3.6, I settle on the more idiosyncratic term ‘extensive conscious mind’.

(8) *Anti-Representationalism*. Representation is a key concept in traditionally-styled cognitive science and features prominently in the orthodox approach to perception following Marr (1982). The sensorimotor theory is unequivocal in rejecting certain misplaced roles for internal representation. These include, in particular, the suggestion that a human perceiver possesses an inner model of the outside world sufficient to generate conscious awareness comprising, all at once, a large, richly detailed or picture-like visual field.

It is unclear whether the sensorimotor theory’s originators intend us to think that sensorimotor knowledge should be thought of as involving representational content. Noë (2004) explicitly declines to take a stance on whether or not cognition in general, including perception, involves or could involve neurally-realised representations. O&N and O’Regan (2011) say they are happy with talk of ‘representation’ in neuroscience when it is used, minimalistically, to describe nothing more than patterns of covariance between physical states. However, they stop short of insisting that there are or could be any representations carrying *content*, which must be evaluable by truth or success conditions, rather than just *information*, which need not, and they stop short of insisting that sensorimotor knowledge *necessarily* involves representation in any sense.

O’Regan (2011), after defending use of the term ‘representation’ in neuroscience, strikes an ambivalent tone by declaring that representation-talk is best avoided. The strongest sign that the sensorimotor approach may necessarily entail the presence of representational content occurs when O’Regan claims that perceptual consciousness depends on sensorimotor engagements plus higher-order thoughts about those sensorimotor engagements. O’Regan states, uncertainly, that the approach might have something in common with classic Higher Order Thought theories of consciousness which, we can note, do typically take it that consciousness depends on higher-order representations of lower-order phenomena (Carruthers, 2011). The sensorimotor theory’s uncertain attitude to internal representation has allowed some commentators to argue, approvingly or disapprovingly, that internal representation plays an essential role in sensorimotor theory as originally presented (Roberts, 2010; Hutto and Myin, 2012). I will later argue that the sensorimotor account, even in the form offered by its progenitors, could in principle dispense with any role in perception for

internal representation.

The conservative variant of the sensorimotor approach consists, largely, in adopting the core claims while rejecting the radical ones. It may also endorse, positively, the following tenets. I will list them quite briefly as they comprise, in essence, the converse of the radical claims:

(9) *Representationalism*. Exercise of sensorimotor knowledge is realised subpersonally by the deployment of internal representations.

(10) *Internalism*. Although perceptual processing is causally scaffolded by bodily exploration, perceptual consciousness is realised solely by the brain.

3 - The Tenets of Sensorimotor Theory

In chapter 2, I identified a number of theoretical posits toward which the sensorimotor theory takes a critical stance, including the bridge locus, pictorial representation, other forms of representation and, more narrowly, phenomena such as filling-in and saccadic suppression. Recall that the sensorimotor theory's critique, properly understood, does not assert that these phenomena could not or definitely do not exist. Indeed, proponents of the sensorimotor theory in some places actively assert that representations and filling-in do play a restricted role in perception. The ambition, as we saw, is to unseat the aprioristic commitment to the idea that perception depends on a bridge locus consisting of a symbolic or pictorial representation, or as O'Regan (1992) puts it, something like an 'inner cinema screen'. The inner screen conception is very close to a view explicitly proposed by Lehar (2003), but, more significantly, is liable to tacitly inform even theorists that do not explicitly endorse it, in particular - as O'Regan observes - those addressing specific phenomena such as filling-in and saccadic suppression. The sensorimotor theory's negative critique is thus similar in spirit to Dennett's (1991) earlier rejection of Cartesian materialism.

Notice that the sensorimotor theory's negative claims do not actively disprove the tenets of the orthodox approach to perception. This does not matter, as long as we appreciate that their ambition is to open up conceptual space for the sensorimotor approach's own theses. As a positive programme, the account draws on a number of unorthodox, ambitious and relatively detailed positive claims about perception and perceptual consciousness, although, as will become apparent, they are frequently in need of a great deal of conceptual clarification. Offering this clarification is my aim in the present chapter, which serves the dual purpose of laying out the tenets of sensorimotor theory and arguing for what I suggest are the most promising ways of understanding them. This will set things up for the remainder of the thesis, in which I defend the sensorimotor theory against some (albeit not all) of the best-known objections, and finish by outlining a contribution the sensorimotor approach could make to a longstanding question concerning perceptual consciousness of temporally-extended events.

3.1 - The World as an Outside Memory

We will see later in the chapter that the relation between Noë's sole-authored version of the sensorimotor theory and orthodox accounts in the tradition of Marr (1982) admits a specific complexity. Although Noë rejects the orthodox subpersonal model, his own account characterises perception primarily at the personal level. This means that a little care is required to see what sort of subpersonal story is needed to do justice to Noë's account, and which subpersonal posits the account invites, rules out or is agnostic about. In its early presentation (O&N), the sensorimotor theory focuses directly on the subpersonal processes that underlie perceptual consciousness, and therefore stands in even more direct rivalry to orthodox subpersonal accounts.

O'Regan (1992) observes that the human retina appears to have a disastrously poor design for retrieving information from the environment. Photoreceptors are increasingly limited in number the further they are from the centre of the retina, meaning that only a small central portion of the retinal image is highly resolved, while the resolution gets increasingly poor toward the outer perimeter. Moreover, the photoreceptors are, throughout the retina, unevenly rather than uniformly dispersed, and are completely absent at the blind spot, an area corresponding to a significant portion of the visual field. Visual perceivers do not ordinarily notice any of these defects, of course, and take themselves to be in perceptual contact with a richly detailed, high-definition visual scene. O'Regan claims that the disparity between the apparent richness of visual experience and the sparse and uneven detail available from the retinal image is a fundamental problem which has not traditionally been addressed with sufficient care. The result, he suggests, is that orthodox theories are frequently under the thrall of a caricatured view of vision in which the defects of the retina are compensated for by the construction of a detailed 3D model of the outside world, the activation of which is necessary and sufficient for perceptual consciousness. Almost no theorists explicitly endorse this conception, but O'Regan claims that it tacitly informs work on a range of specific visual phenomena, for instance saccadic suppression and filling-in. Notice that the conception is not disowned by Marr's very general claim that: "vision is the *process* of discovering from [retinal] images what is present in the world, and where it is" (Marr, 1982, p. 2).

The sensorimotor theory's alternative conception is founded on the claim that conscious visual experience does not comprise a richly detailed representation of the visual scene. A later step will be to re-describe visual phenomenology and the processes that enable it with attention to the fact that perceivers do not take visual experience to be sparsely furnished. The starting point, however, is to establish that the visual system does not process large amounts of detail all at once. O&N, pursuing this thesis, cite a range of empirical work, including work on inattention blindness (Wolfe, 1997), the best known example of which is the 'invisible gorilla' test (Simons and Chabris, 1999). The experiment features a short video of a basketball game with teams dressed respectively in black and white, each passing a separate basketball amongst themselves. Participants are shown the clip and asked to count how many times the players in white pass the ball, a task which requires close attention given the fast-moving action and the distracting influence of the team in black. Midway through the clip, a person in a gorilla costume enters the scene, walks casually through the middle of the game and makes a chest-beating motion. When asked if they had noticed anything out of the ordinary, around half the participants in Simons and Chabris' study failed to report noticing this take place. There are multiple ways to interpret this data, among which is the suggestion that participants were visually aware of the gorilla but did not attend to it, and therefore failed to remember it when questioned. However, the work supports the claim that perceivers do not experience unattended detail from the visual scene.

O&N draw a similar lesson from change blindness, a related phenomenon in which subjects fail to notice sudden changes to the visual scene. In an early study by McConkie and Zola (1979), participants were asked to read a line of text displayed on a monitor in alternating upper and lower case (aLtErNaTiNg CaSe). The experiment was set up so that changes to the text were timed to coincide with saccadic eye movements. During saccades, the text was modified so that the lower case letters became upper case and vice versa, although the wording, typeface, and other features stayed the same. Participants were asked if they noticed anything irregular, and universally failed to report noticing the changing case. This led the experimenters to conclude that 'visual information', meaning information derived directly from the retinal image, is not retained between fixations, although they note that the participants would have had to retain information of a more abstract, non-visual kind given

that they were able to comprehend the meaning of the text, which was read over multiple fixations.

Early work on change blindness (see also, O'Regan, 1984) focused specifically on the consequences of eye movements, and was not motivated by broader theses about the character of vision. Later, O'Regan's (1992) pre-cursor to the sensorimotor theory prompted O'Regan and colleagues to pay renewed attention to the phenomenon (Rensink, O'Regan and Clark, 1997; O'Regan, Rensink and Clark, 1999). Study participants were shown colour photographs on a computer display. Each photograph was alternated with a copy in which a key feature had been modified, for example with an object added or removed, or with the colour of the scenery changed. With the photos switched back and forth, participants could in the ordinary condition easily identify the change. The researchers found that inserting a brief flicker in-between the frames (Rensink et al, 1997) or superimposing an assortment of shapes in a 'mud splash' pattern at the moment of transition (O'Regan et al., 1999) made participants far less likely to identify the change. To explain this, they suggested that perceivers notice visual changes by having their attention drawn to sudden changes in retinal input (known as transients). This means that if the modified feature of the photograph, a local transient, is replaced by a global transient, like the flicker, or supplemented by several competing local transients, like the mud splash, this compromises perceivers' ability to detect the change.

There are various ways this observation might be incorporated into a theory in which vision also depends on the activation of an inner model. For example, the visual system may require transients to update the inner model in line with the changes. Alternatively, the inner model may be sensitive to changes in the visual scene whether or not they are signalled by transients, while transients are required to draw perceivers' attention to these changes. However, the change blindness findings make a uniquely good fit with the claim that there is no inner model, since the finding that transients are a necessary condition for detecting change holds open the more straightforward possibility that transients are sufficient for detecting change even in the absence of an inner model.

O'Regan (1992) offers an alternative framework for understanding vision, the idea of the 'world as an outside memory'. The conception is partly based on an analogy between seeing and remembering. In his examples, you might be reminiscing about your grandmother,

or attempting to remember what you had for breakfast. Long-term memories are not always consciously in mind, but typically come to awareness after being deliberately sought out. When you probe your long-term memory, past episodes do not arrive fully-formed in all their detail. Instead, you engage in a process of asking yourself questions, and answering them, and this way accessing facts about the past in a piecemeal way.

According to the outside memory view, vision works in a similar way, although the role of the memory store is played by the outside world and not the brain. The perceiver makes use of shifts in attention along with bodily movements, especially saccades, to probe the visual environment. This probing involves accessing visual detail serially, a little at a time, and in accordance with the perceiver's or visual system's present interests, i.e. the facts they want to find out and the locations they want to explore. Rather than subserving the construction of an inner model, this activity of temporally-extended probing simply takes the inner model's place.⁹

Another key analogy underlying the outside memory view of vision is touch. While vision is prone to invite an intuitive commitment to the inner model conception, it is very natural to think that perceiving the world by touch is a process of temporally-extended bodily probing rather than model-building. Noë (2004) illustrates this point with the example of a blind person who uses a cane to sample the environment one detail at a time, while making use of an accompanying grasp of the movements required to access further detail. This is how vision works on the outside memory view.

Return, briefly, to the empirical work on inattention blindness and change blindness. The outside memory thesis fits well with the suggestion, derived from work on inattention

⁹ The outside memory thesis says that perception is constituted by temporally-extended probing, and therefore might imply that the substrate of perceptual consciousness is wider than usually thought, comprising for example activity by motor areas of the brain or even activity by the extra-neural body and environment (see section 3.6 and Hurley, 1998). O'Regan (1992) stops short of claiming this, and does not offer any explicit stance on the metaphysical question of how far the substrate of consciousness extends. His early paper actively endorses the view that internal representation plays a role in vision, which might imply that the deployment of the internal representations is sufficient for perceptual consciousness without calling on any strictly necessary extra-neural activity. O'Regan should more realistically be taken as endorsing the milder claim that temporally-extended bodily probing plays an essential role in any explanation of vision.

blindness, that perceivers only experience those features that they attend to. It also predicts that visual systems rely on transients to detect change. The thesis claims that a transient occurring at the periphery of the retinal image causes the perceiver to make appropriate bodily movements to bring the object from which the transient originated into better view. Notice that the retinal image does not need to be highly or uniformly resolved to feature transients, and using transients this way could allow the visual system to detect important events at the periphery of the visual field without using an inner model. With this hypothesis, O'Regan (1992) accurately predicted the outcome of the later work on change blindness, which confirmed that change detection frequently depends on transients.

The outside memory view resonates with Brooks' (1991) suggestion, from situated robotics, that the world need not be internally modelled because it can serve as its own model. In a changing environment, it makes good sense to access detail only when required, since a notional inner model would either need to be constantly updated, which is inefficient, or risk representing the environment inaccurately. The view also observes that the visual system need not compensate for the instability of the retinal image by constructing a stable representation of the outside world. Instead, the visual system simply assumes that the outside world is a persistent and stable entity, and takes advantage of this by accessing detail directly from the environment when required, in a piecemeal and discontinuous way.

Perceivers must, of course, somehow grasp how presently accessed detail forms part of the larger visual scene. O'Regan (1992) claims that retinal stimulation is accompanied by a "non-metric awareness" (p. 474) of the locations of specific objects in the visual scene, which incorporates an understanding of the eye-movements or other bodily movements needed to access more detailed information about them via the retina. O'Regan claims that you may have this kind of awareness even of objects you have your back to, implying that the awareness is neither pictorial in its phenomenal character, nor properly-speaking visual at all. He suggests that it is enabled by neural representations which are symbolic rather than iconic, and sparse in detail. The representations do not themselves realise visual awareness, but help enable the temporally-extended probing that does.

The idea resembles the 'deictic' representation advocated by Ballard and colleagues (Ballard, Hayhoe, Pook and Rao, 1997) in which the visual system does not represent a

detailed environment, but represents sparse facts about the agent's bodily relation to objects in the visual scene. Note that the sensorimotor theory, as it emerges in later work (especially O&N), endorses the idea that perception is a process of temporally-extended exploration assisted by an implicit understanding of the movements required to bring objects into view, but does not commit itself to a particular account of the subpersonal representations that enable this, nor unequivocally to the existence of representations at all.

Return now to the apparent disparity between visual consciousness, which perceivers tend to think presents a richly detailed visual scene, and visual processing, which evidence suggests does not feature detailed representation. The apparent mismatch is captured by the work on inattentional blindness. Subjects in inattentional blindness experiments like the invisible gorilla test typically express surprise when they have the events they missed pointed out to them, indicating that they have certain assumptions about the character of their own experience (Dennett, 2001). This has led theorists like Blackmore (2002) and Dennett (2002) to propose that visual perceivers are in the grip of a 'Grand Illusion' in which they believe that they see a great deal more detail than they actually do. O'Regan (1992) claims that the apparent experience of richness is explained, in part, by the expansiveness of the retinal image, which in spite of its defects, incorporates a relatively large visual field.

This claim is not featured in the sensorimotor theory proper, but O&N endorse O'Regan's further point, which is that perceivers experience richness because they have access to missing detail when required. O&N compare this to the fridge-door effect, in which the refrigerator light seems like it is always on because it is always on when you open the door to look. This is compatible with the claim that perceivers are indeed beset by the Grand Illusion. We will see in section 3.4 that Noë makes a slightly different point, namely that the experience of detail is not an illusion needing to be explained by the fact that perceivers have access to missing detail, but a genuine experience of detail, constituted by perceivers' knowing that they have access to missing detail and implicitly knowing the precise movements that would bring that detail into view.

3.2 - Phenomenal Qualities: The Sensorimotor Contingency Account

A phenomenal quality, such as the look of yellow, the smell of a rose, or the sound of a bell, is a property of perceptual consciousness. In particular, it characterises consciousness from the first-person perspective of the conscious subject. Although we could, as a third party, describe some properties of a person, or their brain, and say we have identified the locus of some specific kind of consciousness, it is an open question whether we could ever thereby be characterising the phenomenal qualities that comprise their conscious experience. Some philosophers, most notably Jackson (1982) and Chalmers (1996), have argued that there is no way, even in principle, to characterise phenomenal qualities using physicalistic vocabulary, and hence that the qualities are not physical. Phenomenal qualities are often referred to as 'qualia', frequently though not exclusively to denote the conception of them under which they are irreducible and non-physical.

3.2.1 - The "Qualia" Problem

Jackson (1982) famously argued that qualia are non-physical using his 'Mary's room' thought experiment. The scenario features a scientist called Mary who is confined from birth to a room where she is deprived of visual stimulation in any colour other than black and white. Jackson asks us to imagine that from inside the room Mary comes to learn literally all the physical facts there are to be known, among which are all the facts that obtain about the physical events that occur when we experience colour. One day, Mary is released from the room and becomes directly acquainted for the first time with a red object. Jackson's position, at the time he formulated the thought experiment, was that Mary now learns a new fact, namely what it is like to see red, and that since she already knew all the physical facts she must therefore have learnt a non-physical fact. From this, he concluded that qualia are not physical. More recently, Chalmers (1996) has used a thought experiment to argue for the broader conclusion that phenomenal consciousness in general is not physical. He suggests that there is a conceivable and therefore, as he takes it, metaphysically possible world which is physically identical in every respect to our own, containing creatures ('zombies') that are physically identical to ourselves but lack qualia entirely. He argues that since the physical properties

instantiated in this world could be instantiated without any attendant phenomenal consciousness, a physical property cannot be identical to a phenomenal quality.

Anti-physicalist arguments like these play quite an important background role even in debates among proponents of physicalism about phenomenal character. Physicalists might deny, for one reason or another, that the premises Jackson and Chalmers use to show that qualia are non-physical are well supported in the first place. For example, some deny that there is any reason to accept Chalmers' premise that whatever is conceivable is possible (see, Gendler and Hawthorne, 2002). Others argue that given the present state of physics, we have no way to tell whether Mary, who knows physical facts as yet unknown to people in the real world, would acquire new knowledge after leaving the black and white room (Montero, 1999). In spite of arguments like these, physicalism is strategically best placed to overcome the intuitions underlying anti-physicalist thought experiments if it can provide compelling positive reasons for thinking one physical property or another is associated, necessarily, with a given phenomenal quality. For this reason, an account of phenomenal qualities has an advantage if it shows better promise than its rivals at overcoming anti-physicalist intuitions, even in a debate conducted exclusively between physicalists.

3.2.2 - Qualia and Neural-Correlates of Consciousness

The threat posed to naturalism by the suggestion that phenomenal qualities are mysterious or even non-physical is a notable and sometimes explicit background presence in the sensorimotor account. O&N begin their canonical early statement of the sensorimotor theory by criticising, in particular, attempts that have been made in neuroscience to account for consciousness by appeal to types of brain activity. They cite a number of neuroscientific accounts that have attempted to explain the existence of consciousness by appeal to facts about the neural activity it is correlated with (neural-correlates of consciousness or NCCs). These facts include, for example, the quantum properties of microtubules (Penrose, 1994), and in another case coherent oscillations within a 40 - 70 Hz frequency range, i.e. potential energy that is repeatedly displaced and restored to its original state at a constant speed of 40 to 70 times per second (Llinás and Ribary, 1993). O&N also target attempts to account for the

differences in phenomenal character between sense modalities, for example vision and audition, by appeal to the distinct neural pathways and cortical regions that are active respectively when we see and hear, a view for which they cite Müller (1838) as the historical precedent, and which they claim most present day scientists continue to be satisfied with. O&N object that all these approaches suffer from the same defect, namely that they fail to give any principled account of *why* the entities, areas or types of activity are correlated with consciousness in general, or with one phenomenal character rather than another. They argue that as long as we persist in appealing to neural properties of this sort an explanatory gap will persist, because there will always remain an unanswered ‘why?’ question.

Even if we set aside worries about anti-physicalism, accounts of the kinds described above have the drawback that they explain less than do accounts that offer a more principled explanation for why one variety of physical state rather than another yields consciousness. O’Regan and Noë’s point is even more obviously fitting when it comes to explaining degrees of difference and similarity of phenomenal character within and between sense modalities, considering that even if you believed that the relation between consciousness in general and the physical properties that instantiate it will never be intelligible to us (see, e.g., McGinn, 1991), it would still be reasonable to hope that a physicalistic account can be given which will intelligibly and non-arbitrarily explain why specific kinds of experience have one phenomenal quality rather than another. The explanatory limitations O&N identify are even more pronounced in light of the threat posed by the anti-physicalist arguments offered by philosophers like Jackson and Chalmers, since the less intelligible the proffered relation between consciousness and the physical, the greater the motivation remaining to believe that physicalism is false.

O&N do not address more promising variants of the NCCs approach. They observe that we cannot read off the properties of consciousness from the topographical quality of cortical maps that are organised in ways that replicate the retinal image, since the retinal images, being dual in number, obscured at the blind spot, and defective in various other ways, do not correlate with what we take to be the felt quality of visual consciousness. O&N do not directly address the possibility that the character of consciousness could be explained by appeal to other topological or topographical isomorphisms, nor the more promising

possibility that phenomenal character could be explained purely by appeal to the content and functional role of symbolic representations. Pessoa et al. (1998) and Noë and Thompson (2004), however, give arguments against these proposals, as I discussed already in chapter 2.

3.2.3 - The Sensorimotor Contingency Account

O&N attempt to improve on these defective explanations of phenomenal character by appealing to *sensorimotor contingencies* (SMCs), the systematic ways that stimulation received to the brain from the sense organs varies in line with movement by the perceiver's body or the objects perceived. One explanatory task that SMCs carry out is accounting for the differences in phenomenal character that hold between sense modalities, for example vision, audition and touch. As an example, O&N propose that you are facing a horizontal straight line. If you face the straight line head-on so that it is perfectly aligned with the equators of your eyes, the image projected onto your retina when flattened appears as a straight line. However, if you move your head or eyes so that the line meets your eyes above or below the equator, the eyes' curvature means that the image projected onto the retina will appear when flattened as an arc. Moving your eyes horizontally, by contrast, leaves the projection of a straight line intact. The appearance of an object also alters when projected at varying locations on the retina in line with the varying character of the photoreceptor cells. The further photoreceptors are from the centre of the retina, the greater their size, but the lower their density. Light meeting different areas of the retina is also met by a differing selection of photoreceptor types. Rods are more sensitive than cones. Cones require more light to become active, and activate, depending on their type, in the presence of differing wavelengths of light. In one part of the retina, the blind spot, there are no photoreceptors at all.

O&N call the patterns of sensorimotor contingency modulated by characteristics of the sense organs *apparatus-related SMCs*. They contrast the apparatus-related SMCs associated with the eyes with those associated with the ears, observing that sense input from the ears is not affected by blinks or eye movements, but does obey its own distinctive patterns of sensorimotor contingency. Tilting your head toward a sound source, in their example, increases the amplitude of the signal but leaves the frequency unchanged. The crucial point

these examples are supposed to illustrate is that vision and sound are not experienced as distinct conscious modalities because of the intrinsic character of the sense input signals or the intrinsic character of the brain regions associated with the respective senses, but solely on account of the distinct apparatus-related SMCs associated with the respective sense organs.

O&N also use SMCs to account for differences in phenomenal character within sense modalities. The range of phenomenal qualities that comprises a particular modality is explained by SMCs determined this time by properties of the objects perceived, rather than the sensory apparatus. These are *object-related SMCs*. Although it is intended to extend to other modalities, the account focuses on vision. Within vision, the most straightforward example of SMCs concerns shape. The sensory stimulation your brain receives from a square-shaped object changes in line with movement by the perceiver or object in a systematic way that differs from the movement-related dependence that holds when you look at a circular object. Unlike colour, the conscious experience of shape does not have any appearance of being irreducibly subjective; simply knowing that a square is a closed shape with exactly four equal sides is in principle sufficient for knowing what a square looks like. The advantage of the sensorimotor account, when it comes to shape, is that it shows how we come to experience shapes as they are, even though they are distorted when they appear on the retina. The account of shape experience is also important because it provides a partial analogue of colour experience.

Colour experience is also accounted for by the systematic dependence between movement and differing sensory stimulation, in particular the activation of differing sets of retinal cones (O&N). The analogy with shape is complicated, however, by the need to account for how the retinal stimulation you receive from an object does not just vary as your spatial alignment with it changes, but also in line with changes in ambient lighting. This point is made more explicit in the longer solo work by Noë (2004) and O'Regan (2011). We therefore need to appeal to three-way patterns of counterfactual dependence between movement, changes in ambient lighting conditions and the stimulation received by the retina. In line with the observations made about colour by O'Regan (2011) - which he points out do not constitute a final or definitive sensorimotor account of colour - this three way relation can be reduced, essentially, to a two-way relation between changes in the character of the light that

meets the object before being reflected onto the retina, and changes in the character of the retinal stimulation that follows. Note, however, that the fact that the two-way relation features fewer variables does not mean it is more appropriate than the three-way relation, which characterises the effects of movement, when it comes to explaining how perceivers actually come to be in states characterised by the relevant phenomenal qualities.

In this way, the sensorimotor account does not just appeal to properties of the light hitting the retina, but to the object's propensity in the presence of varying light sources to reflect varying wavelengths of light, and therefore stimulate the retina in systemically differing ways. O'Regan (2011) argues it is not surprising that this property of the object properly enters into an explanation of colour experience, since we are evolved mainly to respond to objects, and only occasionally interact directly with sources of light. To provide further support for the sensorimotor perspective, he cites some interesting work conducted at his lab. Philipona and O'Regan (2006) looked at changes in how photoreceptors are stimulated in the presence of a range of surfaces and lighting conditions. There are three kinds of cone photoreceptor, responding respectively to light of short, medium and long wavelengths, and each kind can be activated independently of the other two. This means that to describe cone activation in the presence of a given surface, in a given lighting condition, you need to specify a value for each of the three photoreceptors. On this basis, to specify how these change in different lighting conditions, you usually have to specify a value for three variables in each lighting condition.

The study found, however, that for surfaces with certain reflectance profiles, the explanation can be simplified, since the activations of the three photoreceptors do not vary independently in line with changes to lighting. In the case of red, yellow, blue and green surfaces, it was possible to account for changes in retinal stimulation across varied lighting conditions by specifying values in each lighting condition for only one or two variables. Philipona and O'Regan observe that this correlates with anthropological data showing that red, yellow, blue and green are the four most universally adopted colour classifications (Berlin and Kay, 1969). This provides further support, they suggest, for the idea that colour experience does not just depend on how the retina responds to reflected light of different wavelengths, but on how that responsiveness varies given changes to lighting; notice that by

just attending to the way that light of given wavelengths stimulates the retina, we would lack this explanation for the anthropological data. Since the ambient light reflected from the object to the retina varies in line with movements by the perceiver or object, the colour experience depends on a relation between movement and sense input, although to make this relation intelligible we have to also understand how lighting conditions change in line with these movements.

There is some promise, then, that the sensorimotor account of phenomenal quality can account for the experienced feel associated with distinct modalities, and within the visual modality for experienced colour and shape, solely by describing extra-neural characteristics of the perceiver's relation to the environment. This has a methodological benefit, since outward interactions are easier to observe than neural states. More importantly, it provides a more principled explanation than appealing to brain states alone. In part, this is because it appeals to patterns that are self-evidently and necessarily present in the activity of embodied perceivers. It is also because of the systematic framework it provides for explaining degrees of similarity and difference between phenomenal qualities. Noë (2004) argues that the account therefore counts as an improvement on those that suggest qualia amount to an unexplained 'residue' remaining even after we have accounted for an agent's perceptual engagement with the environment.

3.2.4 - SMCs as Counterfactual Conditionals

SMCs, as we have seen, offer a principled way to account for the varying phenomenal character associated with distinct modalities and different percepts within a modality. In particular, they also offer in principle a way to quantify degrees of difference and similarity between qualities. More has to be said, however, about what it is for a perceiver to be in a conscious state characterised by SMCs, and about the role played by the perceiver's internal states in enabling perceptual consciousness. Most versions of the sensorimotor theory complete this part of the story by reference to attunement or practical knowledge pertaining to SMCs. I will examine this notion in detail in section 3.3 and further in chapter 4, but it is worth signposting now the fact that how we understand sensorimotor knowledge or

attunement will be affected by how we resolve a key ambiguity in the notion of sensorimotor contingency.

As they are usually characterised, SMCs are appropriately described by counterfactual conditionals of the form ‘if the subject moved like this, their sense input would change like this’ and ‘if objects in the environment moved like this, the perceiver’s sense input would change like this’. This construction is offered by Noë (2004) and is also frequently in evidence throughout O&N, for example where the authors say: “Seeing an object consists in precisely the knowledge of the relevant sensorimotor contingencies - that is, in being able to exercise one’s mastery of the fact that if, among other things, you make an eye movement, the stimulus will change in the particular way typical of what happens when you move your eyes” (p. 968). On another construction, SMCs comprise the patterns of dependence between sense input and motor output that actually, rather than counterfactually, characterise the perceiver’s ongoing sensorimotor interactions. This alternative notion is found in the summary of the sensorimotor approach offered by Burhmann, Di Paolo and Barandiaran, who claim that: “According to the sensorimotor approach, perception is a form of embodied know-how, *constituted by* [emphasis added] lawful regularities in the sensorimotor flow or in sensorimotor contingencies in an active and situated agent” (2013). Here, the know-how is not knowledge about SMCs, it is knowledge which when exercised helps *constitute* the SMCs. This notion is sometimes evident in O&N, where the authors claim that exercising sensorimotor mastery consists, in the case of a guided missile, in its exercising “action recipes” (p. 945), which map certain sensory stimuli to certain motor outputs. The action recipes, the authors at one point claim, allow the missile to “*make* [emphasis added] lawful changes in sensory stimulation” (p. 945). The apparent role of the action recipes, in this case, is not to constitute knowledge *about* SMCs, but to allow patterns of sensorimotor contingency to be instantiated.

Burhmann et al. offer an explicit account of some distinct notions of SMCs in play in sensorimotor theory, which they propose can be modelled as dynamical systems. The *sensorimotor environment*, as they call it, comprises the complete set of counterfactual conditional-styled SMCs governing a perceiver’s bodily relation to its environment. This is to say that it describes the changes in sense input that would occur if the perceiver made

particular movements. Characterising the sensorimotor environment therefore only depends on the outward characteristics of the perceiver's body and the environment and not on the perceiver's brain. The authors characterise the perceiver's repertoire of actual bodily engagements within a given sensorimotor environment through what they call a *sensorimotor habitat*, which comprises the movements the perceiver makes in line with given sense inputs, and the changes in sense input that follow. As the authors note, the sensorimotor habitat is not just determined by characteristics of the extra-neural body and environment, but also by the agent's internal states.

As we have seen, the original statements of the sensorimotor theory are explicit, although only some of the time, that they mean SMCs to characterise the sensorimotor environment rather than the sensorimotor habitat. This understanding is also evident in the 'radical enactivist' versions of the sensorimotor account offered by Hutto and Myin (2012) and Myin and Degenaar (2014), which reject O&N's appeal to 'knowledge' but agree that perceivers must be 'attuned' to SMCs that exist, presumably, prior to the perceiver acquiring a sensitivity to them. The sensorimotor habitat is likely to prove useful as a way of characterising the perceiver's perceptual attunement to some subset of the SMCs in the sensorimotor environment. However, the sensorimotor habitat conception of SMCs also obscures some of the insights offered by the original statements of the sensorimotor theory. One thing it does is obscure a distinction between distinct directions of fit. SMCs, in their original formulation, describe the changes in sense input that follow changes in motor output, and never the changes in motor output that follow changes in sense input, and we could therefore think of them as having a motor output to sense input direction of fit. Even if we suppose, as if plausible, that the only SMCs relevant to an account of perceptual consciousness are those sometimes instantiated by the perceiver's actual sensorimotor engagements, it would still be true, on the original statements of the sensorimotor theory (O&N; Noë, 2004; O'Regan, 2011) that SMCs by definition have just one direction of fit, while the sensorimotor habitat describes a systematic dependence running bidirectionally between motor output and sense input.

It may turn out that the purported distinction between 'input' and 'output' is arbitrary rather than reflecting a genuine functional distinction, and therefore that the

distinction between directions of fit does not finally matter either. Nonetheless, the sensorimotor theory in its original form is explicit in individuating phenomenal qualities by appeal to SMCs with a motor output to sense input direction of fit. This is compatible with a view where inputs and outputs are instrumentally useful but dispensable parameters for modelling the workings of the wider loop of activity comprising brain, body and environment. It is also compatible with a view under which the input-output distinction is genuine. So as not to pre-judge this question, and others concerning the role played by internal states in the sensorimotor account, I restrict 'sensorimotor contingencies' (SMCs) to describing aspects of the sensorimotor environment. This of course entails that SMCs fail to explain anything about perceptual consciousness unless we say more about the role played by the perceiver and their internal states, and I examine the existing proposals about this later in the chapter. For now, I will continue by outlining the different, perhaps compatible, characterisation of sensorimotor contingency offered by Noë's (2004) solo account.

3.2.5 - Perspectival Properties and Objective Properties

Noë (2004) supplements the SMCs described by O&N with two further kinds of property, which can be thought of as new kinds of SMCs. Instead of focusing on how neural inputs from the sense organs change in line with movement, he focuses on the varying ways that objects appear to conscious perceivers. In particular, he introduces the notion of *P-property* (short for perspectival-property), an idea which is most straightforwardly explained by means of an example. A dinner plate has the 'objective' property of being circular. Viewed at an angle, however, it has the appearance of an ellipse, which gets wider or narrower as the subject faces it at different angles. The elliptical aspect is a P-property, in this case a P-shape. To experience the plate as being objectively circular, the perceiver must be sensitive to the ways that the P-shape would vary in line with movements by the object or perceiver. Thus objective properties can be characterised by reference to an amended notion of sensorimotor contingencies (SMCs). To experience the P-shape, the perceiver must understand their spatial alignment to the object in an egocentric, bodily way. Noë gives examples like knowing which bodily movements would result in your facing the object head-on, and having the ability to

indicate, using a thumb and forefinger, the shape and size of the space taken up by an object in your visual field. Thus P-property experience depends, too, on appreciating the results of possible movements, and therefore, in a way, on sensitivity to a further variety of SMCs. Noë does not say how this account relates to the account of shape perception found in O&N. We can legitimately take it that the accounts run in parallel, however, given that conscious perceivers must somehow be sensitive to SMCs styled along the lines of the earlier account in order to experience P-shapes and objective shapes.

Colour also admits a distinction between P-property and objective property. In Noë's example, you have a wall painted in a solid white. Its being a particular shade of white is an objective property. However, at a given time of day, the wall might be dappled with sunlight in one area, making parts of it appear an especially bright white, and heavily shaded in others, making it appear dark grey. This merely apparent colour is a P-property. To characterise a P-colour, we do not need to explain how retinal stimulation varies with changes to illumination, merely how the light reflected from the object stimulates the retina, and how that stimulation would change if the light met with a differing selection of photoreceptors. Here, Noë does say something explicit about the connection between his present approach and O&N. He suggests that the experience of apparent colour depends on a subpersonal process and that the sensorimotor contingencies involved in apparent colour are inscrutable to conscious perceivers, implying they are SMCs of the sort described by the earlier account rather than conscious SMCs. We experience objective colour - for instance, the colour the wall is actually painted in - by consciously exercising our practical understanding of how the apparent colour would change given changes, amongst other things, in illumination.

This parallels the suggestion given by Philipona and O'Regan (2006), although it is interesting that Noë's account of colour supplements this account by introducing this distinction between conscious and subdoxastic components of colour experience, which track the distinction between P-colour and objective colour. Noë, like O'Regan (2011), admits a final sensorimotor account of colour experience will involve further considerations still. He introduces one further proposal, namely that the objective colour we experience an object to have depends on the objective colour we take objects in its background to have. Experienced objective colour is determined, therefore, by various relations that objects stand in to each

other as well as the relation they stand in to the perceiver. Although the sensorimotor theory does not claim to offer a complete or final account of colour experience, both Noë and O'Regan make explicit their commitment to the view that colour experience is explained by patterns of movement-dependence, whether those patterns apply only to the perceiver and the object, or between the object and other features of the outside environment. Noë claims, moreover, that the capacities we use to discriminate the properties are themselves sensorimotor, a point I return to in section 3.3 and again in chapter 4.

Since the SMCs featured in Noë's account are relations between movement and, for the most part, features of conscious experience, Noë's account is committed to a distinctively controversial claim about visual phenomenology. Martin (2008) objects that the account misdescribes visual phenomenology, since we rarely experience P-properties, such as the elliptical aspect of the plate. He argues that experiencing a P-property involves adopting the 'painterly attitude', a deliberate effort of the sort a painter makes when attempting to faithfully render a scene with the correct perspective; and it is striking that doing this is actually quite difficult. In response, Noë (2008) clarifies his position by noting that you must experience P-properties even when you do not adopt the painterly attitude on account of the fact that you can only be visually aware of an object from some vantage point. Since you cannot view an object without viewing it from a particular perspective, experiencing its objective properties requires you to understand how your present perspective is contingent on your present spatial alignment to the object, and recognising this entails implicitly appreciating the ways the P-property would vary if you or the object moved. Noë claims in the same piece that you can only perceive P-properties themselves if you also experience the relevant objective properties, since you could not experience a P-property as a P-property if you did not understand the way it varies systematically in line with movement. This means that rather than being 'primitive', the ability to experience P-properties arises simultaneously with the ability to experience objective properties.

Noë (2008) also addresses an important point concerning the metaphysical status of P-properties. In at least one respect, P-properties are subject-dependent, since you cannot describe them without reference to a perceiver. Moreover, descriptions of P-properties can be used to describe the character of a subject's perceptual experience. Nonetheless, Noë (2008;

2012) emphasises that they should not be thought of as being akin to sensations, which are commonly understood as properties of consciousness and not properties of the outside world. He claims instead that P-properties inhere in the outside environment as relational properties borne by objects. As he puts it: “looks [...] are not 'mental intermediaries.' They are, rather, aspects of the things we see” (2008, p. 691) and a little later in the same piece: “I am realist about looks. My view is that the elliptical look of the plate from here, no less than its circularity, is a feature of the visible scene before me, a feature of the world itself” (2008, p. 693). It is worth exploring this claim in a little more detail. The claim that looks are properties of the world implies that they exist prior to and independently of the subject’s capacity to experience them. This seems odd, considering that the word ‘look’, applied adjectivally to a feature of the environment, implies the existence of a subject who sees it.

My suggestion is that we reconcile the idea that P-properties are in some important sense dependent on the subject’s perceptual abilities with the idea that they are properties of the world, and not the subject’s mind, by taking it that prior to featuring in a subject’s conscious awareness the P-properties exist as what Shoemaker (1980), following Geach (1969), calls *mere-Cambridge* properties. These are properties that can change without their object’s causal powers changing, and which therefore are trivial. To update an example offered by Shoemaker, we could, for instance, note that everything now existing has the property of being ‘such that Barack Obama is president of the USA’, while noting that this property could for most objects - a piece of fruit, say - change without the object’s causal powers changing. Shoemaker contrasts mere-Cambridge properties with *genuine* properties in order to denote that mere-Cambridge properties, although they exist, are uninteresting. We could contrast the mere-Cambridge cases with a scenario in which Barack Obama makes it the law that for as long as he is President of the USA, it will be a criminal offence to possess apples. In this event, the property of being ‘such that Barack Obama is president of the USA’ will mean that an apple has a causal power, namely the power to cause you to be arrested, where we take it that an object has a causal power to make X happen if its existence increases the probability of X taking place. In this scenario, the property goes from being a mere-Cambridge property to a being genuine property. Let us take it that P-properties are like this. An object might bear some P-colour or P-shape to a person who does not experience it, for example because they are

not looking at it or they cannot see. P-properties might even inhere in objects in relation to possible but non-existent perceivers. The properties would nonetheless exist, and inhere in the world, where this is contrasted with *in the mind of some perceiver*. But the properties only have causal powers, and therefore become *genuine* properties, if there is a conscious perceiver who experiences them.

3.2.6 - Bodiliness, Insubordinateness, Richness and Grabbiness

Sensorimotor contingencies are a cornerstone of the sensorimotor theory. As we have seen, their role is to explain why one sensory modality is qualitatively different from another and, similarly, why perceptual experiences qualitatively differ from one another within sense modalities. However, they do not address the broader question of why we have experiences that feel, qualitatively, like they are world-presenting rather than feeling like some other kind of experience or lacking an experiential quality altogether. The sensorimotor theory answers this question by appeal to four key notions: bodiliness, grabbiness (O'Regan, Myin and Noë, 2004; O&N), and, in later accounts, partial insubordinateness and richness (O'Regan, 2011). The properties, like SMCs, govern perceivers' sensorimotor activity in a structured environment, which is to say they help characterise the sensorimotor habitat. From a theoretical perspective, as well as from the perspective of the perceiver herself, they offer a principled way to distinguish neural signals related to perception from other kinds of signal.

Bodiliness is the tendency of body movements instigated by the perceiver to bring about changes to the input received from the sense organs. For instance, blinks and eye movements bring about radical changes to the nature of retinal input, head movements bring about changes to the amplitude of the auditory signal received from each ear, and so forth. The fact that the conscious perceiver has a degree of agential control over the character of the sensory stimulation received is supposed to help explain the sense of 'mineness' that accompanies a conscious perceptual experience. It is a necessary component of perceptual experience.

Insubordinateness is the opposite. It describes the tendency for perceivers only to have partial control over the quality of the sensory stimulation they receive. If they had full control,

the perceiver or their visual system would have no reason to think that the sensory stimulation originates from outside themselves. Thus, insubordinateness is also a necessary feature of perception, and helps distinguish neural activity involved in perception from other sorts of neural activity.

Richness is another feature that differentiates perceiving from thinking and imagining. The amount of information accessed in the course of a temporally-extended act of perception is much greater than the information available to thought or imagination. In O'Regan's (2011) example, you might consciously reconstruct, in an act of imagination, the appearance of your grandmother's face. But without seeing her, there are likely to be numerous small details about her appearance, e.g. the width of her eyebrows or the appearance of her hair line, that you cannot reconstruct. O'Regan observes that richness is not sufficient to distinguish perception from every other kind of bodily activity, since autonomic functions such as digestion and respiration also involve rich information about blood sugar levels, oxygen levels, and so forth. However, richness distinguishes sensory information from information playing a similar role in acts of thought or imagination.

Grabiness refers to the tendency of perceivers to shift their attention and bodily alignment in response to transients, sudden changes to sensory stimulation. The fact that perceivers detect change this way makes a vital contribution to the phenomenology of perception, namely the experience of being reliably in touch with a rich and dynamically-changing environment. Grabiness is a key feature of the idea of the world as an outside memory.

The above four properties, for the most part, characterise a perceiver's sensorimotor interactions in a way that can be observed from the outside. For instance, we do not need to know anything about brains or the subjective character of conscious experience to know that retinal stimulation is governed by bodiliness and insubordinateness. We can tell that richness exists purely by virtue of the rich detail available in the environment for pick-up by the retina. Grabiness is supposed to consist largely of bodily movements, in particular eye movements directed toward the source of the transients. This is complicated a little by the fact that grabiness is also supposed to be signalled by shifts in attention, which may not require eye-movements, and may instead depend solely on some kind of neural activity. With this

small complication aside, the four properties, alongside SMCs, help make the phenomenal character of perceptual experience naturalistically explicable without reference to apparently mysterious neural-correlates of consciousness.

However, perceptual experience depends on something more than SMCs and the four other principles, since it is easy to imagine objects standing in relations to the environment governed by these or similar laws without being conscious and even without demonstrating the kind of sensitivity that might justify the suggestion that they are unconscious perceivers. Perception requires that subjects are, somehow, skillfully sensitive to the fact that their relation to the environment is governed by the relevant laws. This is the purpose of sensorimotor mastery, which I discuss next.

3.3 - Sensorimotor Mastery

SMCs, on the usual understanding, are properties that could be described by counterfactual conditionals, and which pre-exist any bodily engagement with the environment, or any capacity for the same. On this basis, it is evident that to explain how a subject comes to perceive and thereby undergo consciousness characterised by the relevant phenomenal qualities, the sensorimotor account needs to show how the subject is, in a relevant way, sensitive to those SMCs. The sensitivity in question is said, in various places, to involve sensorimotor ‘skill’; ‘attunement’ to SMCs; or ‘practical knowledge’, ‘understanding’ or ‘tracking’ of SMCs. Some of the time, the variation in terminology is purely stylistic, but at other times the choice of term reflects some quite different ways of understanding the capacities in question. To keep things as clear as possible, I am going to introduce my own convention and use ‘sensorimotor mastery’ as a broad label to describe any variety of sensitivity to SMCs of alleged interest in sensorimotor theory. Under this broad heading, I will group the capacities in question under three apparently distinct types, which I will call (i) *Sensorimotor Attunement*, (ii) *Sensorimotor Tracking* and (iii) *Sensorimotor Understanding*. In appealing to these three categories I am not prejudging whether there are capacities that fall under more than one of these categories, nor whether all three deserve to make an appearance in the account.

3.3.1 - Sensorimotor Attunement

O&N suggest that over the course of your life, you are exposed to numerous SMCs, and that your brain stores information about them, categorising them under laws of sensorimotor contingency corresponding to various types of property, for example specific colours and shapes. When you are confronted with a red object, for example, your brain deploys information about the SMCs associated with red objects derived from the previous encounters. O&N do not explain what causes the brain, at a particular moment, to deploy stored information about one sensorimotor law rather than another. We can plausibly suppose that the brain does not just make use of present sense input state to detect which SMCs obtain, but also efferent copy from the motor regions, and the way both have changed in the last few moments.¹⁰ This is to suggest that the brain responds to information about the restricted set of SMCs evident in the bodily interactions that are taking place right now, and associates these SMCs with stored information about the fuller set of SMCs that usually apply in similar situations. I think it plausible that this picture, or something close to it, is what O&N have in mind, although they do not say so explicitly.

The authors are emphatic that neither the availability of the stored information nor its mere ‘activation’ is sufficient for perception. Instead, they argue that perception is a kind of embodied exploration of the outside environment analogous to a dance, an activity that can never consist of an inner state alone, but also requires a variety of temporally-extended interaction with the outside environment. They say: “Seeing is not constituted by the activation of neural representations. Exactly the same neural state can underlie different experiences, just as the same body position can be part of different dances” (O&N, p. 966). This invites an understanding of sensorimotor mastery as a capacity for making a variety of systematic, goal-directed bodily movements. This capacity, which I call Sensorimotor Attunement, can also be found in Burhmann, Di Paolo and Barandiaran’s (2013) model, which invites us to model a perceiver’s repertoire of bodily engagements in a structured

¹⁰ It is not a logical necessity, however, that the neural system which enables conscious vision is always sensitive to ongoing efferent copy. It might only be sensitive to efferent copy during learning phases, as Clark (2008) suggests, after which it uses sense inputs alone to recognise patterns of sensorimotor contingency.

environment as a sensorimotor habitat.

Although O&N and Noë (2004) on one or two occasions deny that vision is exhaustively accounted for by a set of behavioural capacities, the idea that vision depends on being poised to act is recurrent throughout their accounts, for example their claim that: “Under the present theory, visual experience does not arise because an internal representation of the world is activated in some brain area. On the contrary, visual experience is a mode of activity involving practical knowledge about currently possible behaviors and associated sensory consequences. Visual experience rests on know-how, the possession of skills” (O&N, p. 846), and in Noë’s account: “It is my possession of basic sensorimotor skills (which include the abilities to move and point and the dispositions to respond by turning and ducking, and the like) that enables my experience to acquire *visual* content” (Noë, 2004, p. 90).

The idea of Sensorimotor Attunement as a behavioural capacity is evident, in particular, in the main example given by O&N, which features a guided missile. The missile’s goal is to hit its target and, while following the target, to try to avoid enemy fire by zig-zagging left and right. The missile’s relation to the target is characterised by various SMCs. When the missile turns left (on its yaw axis), the target, if it appears in its sensors, will appear further to the right-hand side, and vice versa. If the missile gets closer to the target, the target will occupy more space in its sensors. To achieve its goal, the missile’s behaviour must exhibit sensitivity to SMCs like these. O&N, who are emphatic that they do not wish to underplay the contribution made by the brain to perception, propose that the missile, like a human, makes use of stored information, which in the missile’s case comprises “action recipes” (O&N, p. 945) corresponding to various sorts of target. They emphasise, however, that the internal states do nothing until they are deployed for some practical end. In the case of the missile, they observe that the action recipes are “latent” (O&N, p. 945) until the missile begins following a relevant target.

The authors do not say exactly what an action recipe is, but evidently it comprises instructions about how to move. Presumably, it does this by mapping sense inputs to motor outputs, for example specifying that the missile should turn right if the target appears on the right hand side of its sensors. Attending to the contrasting descriptions of the stored information deployed by humans and the guided-missile when exercising Sensorimotor

Attunement, we might notice that direction of fit, which has already been an important notion in understanding SMCs, is an issue again here, although in a different form. The stored information about sensorimotor laws that O&N allege to be involved in human perception appears to have a belief-like direction of fit - it specifies how sense input would change given particular movements, and it must therefore correspond to pre-existing facts about the sensorimotor environment. Action recipes, on the other hand, appear to have a desire-like direction of fit - they specify which movements will take place given particular sense inputs, suggesting that if they are effective it is the missile's behaviour that will appropriately correspond. The fact that O&N refer interchangeably to mastery of SMCs and the possession of 'action recipes' has the potential to be puzzling if we do not unpack the notion of Sensorimotor Attunement with care, since it is not clear straight away why we should think of the belief-like and desire-like capacities as being identical. To this end, we must attend carefully to O&N's claim that the exercise of Sensorimotor Attunement is never *solely* a matter of stored information being activated by the presence of relevant stimuli, but a matter of the stored information being activated in the course of the agent's goal-directed behaviour.

O&N insist that the attunement, is "a practical, not a propositional form of knowledge" (p. 944). Given their characterisation of sensorimotor knowledge as knowledge *that* particular SMCs obtain, we might be drawn to suppose that despite not being propositional, it nonetheless has a belief-like direction of fit. On this basis, their insistence that it must be exercised might reflect nothing more than the commonplace distinction between possession and exercise that applies to any kind of knowledge. For example, your knowledge that Edinburgh is the capital of Scotland only features in your awareness if you are exercising it, perhaps because you are presently thinking about travelling to the capital of Scotland. This offers one reason for thinking sensorimotor knowledge or attunement must be exercised. O&N endorse a stronger claim, however, when they say: "the notion of being tuned, or having [sensorimotor] mastery, *only makes sense* [emphasis added] within the context of the behaviour and purpose of the system or individual in its habitual setting [...] if exactly the same system was being used for a different purpose [...] it might well be necessary for the system to have a different behavior [in order to exercise sensorimotor mastery]" (p. 943). O&N cannot simply be claiming here that a system *ought* to behave differently if its purpose

is different, since this would be trivial. Instead, the passage suggests that the attunement to SMCs is only something that exists by virtue of the agent's purpose and behavioural repertoire.

If Sensorimotor Attunement can only be made intelligible by appeal to a capacity for goal-directed behaviour, as the above quote appears to suggest, it depends on the existence of a capacity with a desire-like direction of fit. In this case, it is reasonable to ask where the belief-like attunement to SMCs comes from. One way to unpack O&N's point is by supposing that we specify which SMCs the perceiver is attuned to by interpreting its behaviour in light of the SMCs that now obtain, along with what we take to be the perceiver's goal state. This might involve something analogous to Dennett's (1987) proposal that we properly ascribe beliefs and desires to a subject on the basis of their explanatory value in making sense of the behaviour. In the case of Sensorimotor Attunement, we first might notice that the missile is disposed to yaw to the right when target appears in the right hand side of its sensors. Next, we might suppose that the missile's purpose is to keep the target aligned in the middle of its sensors; this ascription of purpose could either be a theoretically useful assumption, as with Dennett's model, or we might actually know that this is its purpose because we know what the missile's designer intended. With this purpose in mind, we can make sense of the missile's behaviour by ascribing to the missile attunement to the relevant SMCs, in this case the fact that turning right will result in the target appearing in the middle of its sensors. The 'action recipes' may work by mapping possible sense inputs to appropriate motor outputs, while the action recipe selected may co-vary with different sets of SMCs. It is the interpretive stance, however, that on this proposal admits the dual directions of fit.

O&N do not go so far as to overtly insist that 'purpose' is necessary to specify which SMCs the perceiver is attuned to, either by using the interpretive strategy sketched above or any other strategy. Indeed, neither teleology nor having a proximal goal are given a prominent theoretical role either in the remainder of the paper or in other statements of the sensorimotor account. It is probable, then, that O&N merely mean to suggest that having a purpose is necessary for Sensorimotor Attunement to exist in general. On this basis, their view can be taken as the more austere claim that the perceiver is attuned to a set of SMCs if it is able to respond with a distinctive set of movements when those SMCs obtain, and different

movements when other SMCs obtain, with the added condition that for this sensitivity to constitute *attunement* (or knowledge, skill, etc.) it must somehow help the perceiver fulfil a purpose or achieve a goal. This latter proviso seems right, considering, for example, that when a lake reflects the night sky it demonstrates a sensitivity to the night sky, but does not thereby manifest any skill.

It seems viable, then, that Sensorimotor Attunement could be regarded as nothing more than a behavioural capacity of the sort I have just described. O&N appear to disavow such a view, however, when they make an effort to rebut charges that their theory amounts to a form of behaviourism. In particular, they say: “[I]t is *not* the case that mental states (experiential states) are logical constructions of actual and possible behaviour states” (2001b, p. 1014), and later: “Knowledge of the laws of sensorimotor contingency themselves must surely be represented” (2001b, p. 1017). We could attempt to do justice to this with a representationalist gloss in which Sensorimotor Attunement involves the deployment of representations with both belief-like and desire-like directions of fit, i.e. representations of the SMCs that now obtain and the behaviour that should follow. The notion of a neural representation with both directions of fit is found in the literature on *action-oriented representation* (e.g., Wheeler, 2005; Clark, 1997), and if we wanted to give explanatory pride of place to the neural state that realises the behavioural capacity, it would on this view be the content of the representations that determines the SMCs to which the perceiver is now attuned. However, it is not clear that this fits with O&N’s claim that: “Seeing is not constituted by the activation of neural representations. Exactly the same neural state can underlie different experiences, just as the same body position can be part of different dances” (p. 966).

I think the best way to understand Sensorimotor Attunement is to suppose that while other varieties of sensorimotor mastery involved in perceptual consciousness may draw on something other than a behavioural capacity, Sensorimotor Attunement is, in fact, a logical construction of actual and possible behaviour states. We can do justice to O&N’s resistance to appealing to nothing but behaviour by adding a couple of provisos to the present notion of Sensorimotor Attunement. A first is that a full explanation of the faculty must also explain, by appeal to stored information, representation, or something else, how the behavioural capacity

is realised. The present notion of Sensorimotor Attunement denies that these internal states can be explanatorily relevant if they do not explain the behavioural capacity, and further denies that the perceptual state can change without the behavioural capacities changing. Whether or not the notion amounts on this view to a form of behaviourism is open to debate, but it certainly has the resources to refute the suggestion that because it is a form of behaviourism it ignores or denies the important role played by the brain.

The second proviso is that the behavioural capacities constituting Sensorimotor Attunement cannot be accounted for by a simple mapping from varieties of sense input to varieties of motor output, but only by a mapping from laws of sensorimotor contingency to varieties of motor output. O&N suggest this entails that because sensorimotor laws vary in line with differences in the body's morphology, minute differences in embodiment between one perceiver and another entails that the two perceivers must be discriminating a different range of sensorimotor laws. I will argue later that the sensorimotor theory need not be committed to the view that perceptual states are maximally sensitive to the details of embodiment, although it ought to say that differences in embodiment between two perceivers necessitates that they are making a different range of perceptual discriminations when the differences are big enough.

Going further beyond the explicit claims made by O&N, I am going to add one further criterion for establishing that a perceiver exhibits a belief-like attunement to SMCs of the sort we require. Wheeler (2005) describes work by Harvey, Husbands and Cliff (1994) featuring a robot, designed using artificial evolution, which was able to achieve a practical end by following an identical behavioural repertoire for all sense input states other than the target state. Placed in an environment containing a triangle and a rectangle affixed to a wall, the robot's task was to locate the triangle and move toward it. Its strategy was to do nothing but rotate until its sensors were stimulated in a way characteristic of the triangle, and then move forward. Faced at an angle, the rectangle's edge would appear in the robot's sensors as a diagonal line, and in cases where the rectangle's perspectival shape was sufficiently similar to the triangle's perspectival shape when faced head-on, the robot would 'mistake' the rectangle for the triangle and begin to move forward. However, once the robot had moved some way toward the shape, the rectangle would stop presenting the triangle-like aspect, causing the

robot to resume rotating and eventually head toward the triangle, which would continue to present the appropriate perspectival shape as the robot got closer.

We can observe that the robot certainly possessed a rudimentary degree of Sensorimotor Attunement. It was implicitly attuned to the fact that continuing to rotate while it had the triangle in its sights would result in its losing sight of the triangle, and to the fact that proceeding straight on toward the triangle would allow it to keep the triangle in view. More interestingly, its behaviour over a short time period also manifested a sensitivity to the fact that although the rectangle had the propensity to appear like the triangle when faced from certain positions, it did not from others. This latter sensitivity is not explained by stored information or historical facts about the SMCs that the robot had been exposed to, but by properties of the rectangle itself. Wheeler (2005), following Wheeler and Clark (1999), describes this sort of environmental contribution as *non-trivial causal spread*, denoting that the environment itself plays some explanatorily interesting causal role in the production of the behaviour.

O&N do not say anything explicit to indicate whether temporally-extended engagements demonstrating sensitivity to SMCs, where that sensitivity is accounted for by non-trivial causal spread, should be counted as instances in which the agent exercises Sensorimotor Attunement to the same SMCs. The most cohesive way to understand the sensorimotor account is to suppose that they should not. This reflects the fact that O&N actively emphasise the role played by internally stored information derived from previous encounters with the environment. Indeed, sensitivity to SMCs accounted for by non-trivial causal spread cannot make any contribution to O&N's view that when we exercise the right kind of sensorimotor mastery we are conscious of SMCs that are not capturable by the patterns of output-input dependence manifest in our present movements. For example, you can be conscious of a wall's objective colour without any changes in illumination taking place; and you can feel embedded within a richly detailed visual environment without taking in all that detail, even in the course of a temporally-extended sensorimotor engagement, a point I will discuss in more detail later in the chapter.

Since non-trivial causal spread cannot account for many of the key phenomena the sensorimotor theory is supposed to account for, I think it best to consider the non-trivial

causal spread case to demonstrate that relatively basic Sensorimotor Attunement *plus* non-trivial causal spread can result in surprisingly sophisticated patterns of sensorimotor engagement, but avoid attributing those patterns to Sensorimotor Attunement alone. Another way to put the same point is to stipulate that possession of Sensorimotor Attunement cannot depend on anything temporally-extended, even if it is exercised in the course of temporally-extended behaviour. To be in the relevant way attuned to SMCs that now obtain, a perceiver must have the ability to make an immediate change to its motor outputs demonstrating its sensitivity to those SMCs.

Sensorimotor Attunement, the way I have characterised it, reflects a notion found at least sometimes in the main statements of the sensorimotor theory. It is also, in principle, compatible with Myin and Degenaar's (2014) notion of 'attunement' to SMCs, although it is worth noting that their notion is distinct. Sensorimotor Attunement, as I have presented it here, implies a capacity for outwardly observable behaviour, with representation playing no more than an instrumental role in enabling this behaviour. Myin and Degenaar also describe attunement to SMCs as an embodied capacity, but say: "in the sense of 'embodied' at stake here, 'embodied' is used in contrast with 'representational'" (p. 1). By this, they mean to contrast their preferred account with those accounts that mention embodiment *and* representation. However, Myin and Degenaar's preferred understanding of 'embodiment' also stands in contrast with accounts that make explicit appeal to capacities for outward behaviour.

They stop short of making any such appeal to any such capacities, referring only to the agent's *historical* sensorimotor interactions. They say: "Attunement to sensorimotor contingencies then means that an organism has acquired, on the basis of a history of interactions, a sensitivity in its perception and action for the ways stimuli change with movement" (p. 5). But other than denying that it involves representation, they do not say anything about what a sensitivity to SMCs in perception is constituted by right now. Note, then, that although Sensorimotor Attunement as I have presented it does not rule out representation, it provides a positive non-representational account of what it is to possess the relevant capacity, namely the disposition to manifest different behaviour when different SMCs obtain. This means the notion I outline now has an explanatory advantage compared to Myin and Degenaar's, which is limited in that it only tells us what caused the organism to be

in its current state, and not what it is about the organism's current state that results in it being in a perceptual state.

3.3.2 - Sensorimotor Tracking

O&N note the existence of an apparently key limitation with the notion of Sensorimotor Attunement as I characterised it in the last section - namely that since the guided missile, for example, is not conscious, Sensorimotor Attunement is not sufficient for perceptual consciousness. To account for visual awareness, O&N introduce the further criterion that you must be exercising your attunement to SMCs for the purpose of "thought and planning" (p. 944). To illustrate, they give the example of driving a car while talking to a friend. While driving, you must navigate a range of obstacles while modulating your driving behaviour appropriately, yet you are not necessarily aware of everything you are doing. In this respect, you are equivalent to the autopilot of a plane. Only when you attend, for example, to the colour of the car in front, and think or talk about it, or make plans in accordance with it, do you become aware of the colour. One way to understand this requirement is to suppose that the Sensorimotor Attunement only yields perceptual consciousness when it is deployed in certain kinds of bodily movement. Another possibility is that it depends on a certain kind of neural representation; O&N do not give a particularly detailed proposal about what happens physically when thought and planning takes place. I will call sensorimotor mastery in this context Sensorimotor Tracking.

In his book *Why Red Doesn't Sound Like a Bell*, O'Regan (2011) builds on the account offered by O&N. The book explains in greater detail the relation between Sensorimotor Attunement, which by itself does not necessarily involve consciousness, and the thought and planning needed for conscious perceptual awareness. The key faculties described in O'Regan's model of perception and perceptual consciousness comprise four hierarchically-organised tiers. The bottom tier - call this Tier 1 - comprises the agent's skill-driven sensorimotor coupling with the outside environment. This type of coupling could be instantiated by a human, but also by an unconscious artifact like the guided missile.

O'Regan's account of sensorimotor mastery and SMCs at Tier 1 at times appears to

adopt the notion of SMCs in which they are constituted by patterns present in the agent's repertoire of skillful behaviour (its sensorimotor habitat), rather than the notion in which the skillful behaviour involves a behavioural sensitivity to counterfactual conditional-styled SMCs that pre-exist in its sensorimotor environment. As I have already argued, it is better to explicitly reject this understanding of SMCs, for reasons of expediency. The understanding obscures the distinction (which *may* be theoretically useful) between patterns of dependence between sense input and motor output, and those between motor output and sense input. It also obscures the distinction between SMCs which are merely manifest in this behaviour, and those to which the agent's behaviour demonstrates a skillful sensitivity. Finally, it muddies our understanding of SMCs in light of the contradictory conception on them found in Noë (2004), throughout much of O&N, and at times in O'Regan (2011), in which they are the ways that sense input *would* change given particular movements, regardless of whether those movements take place. Fortunately, it is also consistent with O'Regan's (2011) account if we regard Tier 1 as involving unconscious sensorimotor coupling drawing on Sensorimotor Attunement to SMCs as I have characterised the concepts in the last sections.

The next tier in O'Regan's hierarchy - call this Tier 2 - involves a higher-order cognitive access to SMCs governing activity at Tier 1. This is Sensorimotor Tracking, and is exercised when the subject is poised to exhibit sensitivity to sensorimotor laws in the course of "planning, decisions, or other rational behaviour" (O'Regan, 2011, p. 121). The behaviour in question falls short of self-conscious behaviour or discursive thought and appears to be quite intimately connected, still, with bodily behaviour; O'Regan offers as an example of behaviour at this tier the state of being poised to press on a brake pedal when a red traffic light is visible. It is not clear whether or not we should suppose that the only SMCs you can be sensitive to in Tier 2 activity are those you already exhibit Sensorimotor Attunement toward in your unconscious activity. However, O'Regan's account does not make this an obvious requirement. So it is more straightforward to suppose that Sensorimotor Tracking at Tier 2 involves exhibiting sensitivity to SMCs, in general, for the purpose of behaviour which is planned and answerable to rational norms. This depends on tracking or being sensitive to the SMCs governing your unconscious activity at Tier 1. But it may involve being sensitive to SMCs beyond those featured in your Tier 1 Sensorimotor Attunement, considering that some

of the patterns manifest in your sensorimotor activity at Tier 1 may not be SMCs to which this activity is itself Attuned, considering, to take one example, the proviso I made in the last section that Sensorimotor Attunement does not encompass skillful sensitivity to SMCs when this is accounted for by non-trivial causal spread.

O'Regan suggests that Tier 2 activity involves the sort of cognitive access to some feature that Block (1990a) would describe as having 'access consciousness', since it involves being poised to make use of that feature for reasoning and planning. O'Regan denies that having this first order of cognitive access is sufficient to constitute being conscious, however. To be conscious, O'Regan suggests we must have a further order of cognitive access to activity at Tier 2. This further higher-order tracking, call it Tier 3, constitutes another level of Sensorimotor Tracking. Perceptual consciousness occurs when the perceiver thinks reflectively about features of the perceptual environment corresponding to SMCs which it already tracks at Tier 2. The subject does not have access, when engaging in this kind of thought, to the individual SMCs that correspond to a particular property, but merely to categories or descriptors corresponding to sensorimotor laws, e.g. 'that shade of red'. This kind of thought is not intimately connected with bodily activity. However, in common with the activity described by the tiers below, it is always practically focused, involving active, goal-directed interrogation of the visual scene. In O'Regan's example, this involves asking and answering questions like: "Is it an 'A'?; 'Is it Times New Roman?'; 'Is it black on white?'" (2011, p. 28).

In spite of the more fully-developed character of O'Regan's (2011) account, it lacks a detailed characterisation of the neural (or neurally-realised, functional) processes that enable activity at Tiers 2 and 3. At one point O'Regan suggests his account may bear an affinity to Higher Order Thought (HOT) theories of consciousness, although he is explicit about not being firmly committed to this suggestion. HOT theories explain consciousness by appeal to higher-order access to lower order states in a manner similar to O'Regan's model. Notably, they almost always cash this suggestion out by appeal to a subpersonal explanatory level in which one part of the cognitive systems bears higher-order representations of activity in another system. If the tiered sensorimotor theory adopted a view like this, it would be committed to explaining Tier 3 by reference to a subpersonal representation of Tier 2 activity,

and Tier 2 by reference to subpersonal representations of SMCs. This sort of proposal will be particularly relevant in chapters 5 and 6, where I discuss the relation between sensorimotor theory and the dual systems model of vision.

O'Regan (2012) later asserts that the tracking carried out by Tiers 2 and 3 involves neural states that 'track' or 'register' the appropriate lower-order activity, but do not constitute representations. The neural states register the appropriate interactions by virtue of being the sorts of physical states that ordinarily occur when the appropriate interactions are taking place. O'Regan denies that these states thereby count as representations, although it is notable that some accounts of representation (e.g., Tye, 2000, or Rupert, 2011) appeal to no more than simple covariance relations of this sort to constitute representation, as I will discuss in chapter 5. As it is, O'Regan's (2012) notion of tracking appears to correspond to Myin and Degenaar's (2014) notion of attunement to sensorimotor contingencies (not to be confused with 'Sensorimotor Attunement' as I use it here) in which the agent is attuned to sensorimotor contingencies by virtue of being in some unspecified physical state while having the relevant history of sensorimotor interactions.

The notion of tracking or registering just described appears to be inadequate by itself to do justice to O'Regan's (2011) multi-tiered account, however, inasmuch as it does not explicitly specify the role these physical states perform when deployed in reasoned behaviour (at Tier 2) or reflective thought (at Tier 3). Considering these activities are supposed to be essential for enabling perceptual consciousness, it is reasonable to desire a better developed, naturalistically-respectable account of what these higher order capacities involve. One possibility is that Tier 2 and perhaps even Tier 3 capacities depend only on capacities for behaviour, for example behaviour that is suitably sophisticated, or answers to appropriate rational norms. O'Regan does not say much about this, but suggests possible sympathy toward the suggestion that we ascribe these capacities to an agent through a Dennett-style interpretive stance similar to the one I discussed in Section 3.3.1. O'Regan hints at a position of this nature in a footnote, where he declares: "The implication is that conscious access is not something in the agent. It is a way that an outside person (and consequently the agent itself) can *describe* the state of the agent" (2011, p. 125n1).

3.3.3 - Sensorimotor Understanding

Sensorimotor Understanding is the variety of sensorimotor mastery featured in Noë's (2004; 2012) sole-author versions of the sensorimotor account. I examine the capacity in greater detail in the next chapter, so for now I will restrict myself to a brief outline of its key features. In contrast with the other varieties of sensorimotor mastery, Sensorimotor Understanding pre-supposes consciousness. For this reason, Noë does not present his version of the sensorimotor account as an attempt to dissolve the hard problem of consciousness (cf., O&N; O'Regan, 2011) and Sensorimotor Understanding is instead merely intended to help characterise perception and perceptual consciousness in an already conscious perceiver. Noë (2004) describes the capacity as an implicit understanding of P-properties, the understanding of which requires an appreciation of your egocentrically specified alignment to objects in the environment, and objective properties, the understanding of which requires an appreciation of how P-properties would change in line with movement. Although he downplays the importance of the personal/subpersonal distinction, Noë's (2004) characterisation of Sensorimotor Understanding presents it primarily as a personal-level phenomenon. However, he suggests that for some kinds of experience, including P-colour experience, we do not have conscious access to our exercise of sensorimotor mastery, thereby implying that in limited cases it is firmly subpersonal.

Noë and O'Regan's respective sole-author accounts offer quite different treatments of sensorimotor mastery and there is little to suggest that Sensorimotor Tracking and Sensorimotor Understanding should necessarily be regarded as different descriptions of the same notional faculty. However, given that they both originate from the early O'Regan and Noë paper, they could, with a little finessing, be thought of as different treatments of the same phenomenon. If so, Sensorimotor Understanding could be regarded as a more philosophical exposition of Sensorimotor Tracking. The main obstacle to such a reconciliation is the difference between the properties the respective accounts claim are the objects of sensorimotor mastery, namely subdoxastic sense inputs (O&N; O'Regan, 2011) and P-properties and objective properties (Noë, 2004; 2012). Noë's (2004) account does make reference to the patterns of movement-related change governing input from the sense organs,

claiming we must master them to experience P-colour. This indicates that Noë does not intend his sole-author account to entirely supersede the early statement of the sensorimotor theory, and thereby makes a reconciliation between his and O'Regan's later solo accounts look more promising. Sensorimotor Tracking and Sensorimotor Understanding notably both lack a proper account of the physical or subpersonal mechanisms they are realised or constituted by.

Sensorimotor Understanding is uncertainly defined in a couple of key respects. One, as I just suggested, concerns the capacity's place in a scientific account of perception, in particular whether or not it depends on internal representation, and more narrowly the precise representational or non-representational mechanism that enables it. Another concerns Noë's (2004) claim that it is a kind of knowing-how, and not a kind of knowing-that, which is in tension with his firm insistence, elsewhere in the book, that Sensorimotor Understanding is not knowing how to act but knowing what would happen if you did act. I make some suggestions about how we should resolve these issues in chapter 4.

3.4 - Perceptual Presence: As-Access and In-Absence

As we have seen, the sensorimotor theory frequently criticises the idea that vision is enabled subpersonally by the neural activation of a detailed or, in particular, pictorial representation of the outside world. The outside memory thesis says that the brain does not internally model the outside environment but instead accesses information directly from the outside world. This entails that perception, considered subpersonally, depends non-trivially on bodily interaction with the outside environment, and therefore does not depend solely or primarily on internal representation.

In his sole-authored versions of the sensorimotor theory, Noë turns his attention to a personal-level analogue of the subpersonal views targeted by O&N. The 'snapshot conception' (Noë, 2004) endorses a pictorial understanding of the phenomenology or what-it-is-likeness of visual experience, suggesting that visual experience presents in pictorial fashion a large and uniformly detailed visual field. Noë (2012) draws a connection between this phenomenological claim and the logical claim that perception at the personal level involves

representing the outside world. Declaring an allegiance with relational realism (e.g., McDowell, 1994a; Campbell, 2002), he denies that conscious perception involves personal-level representation, and claims instead that it involves standing in a specific kind of relation to the outside world itself. This point is reflected by the word he uses to label the objects and properties to which we are intentionally-directed in conscious perception, which he describes as being 'present' rather than 'represented'.

Noë claims that objects and properties in the outside environment become perceptually present when we stand in a particular kind of bodily relation to them. As we have seen, that bodily relation can be specified by counterfactual conditionals describing the effects of possible movements, with the relevant counterfactual conditionals being ones the perceiver implicitly understands. Noë's (2012) relational view of intentionality generalises to encompass non-perceptual forms of thought, with perceptual presence being distinguished from other forms of presence by the fact that the perceptual relation belongs to this particular type. Crucially, the presence perceivers achieve by standing in this relation is not the simple kind of presence characteristic of the snapshot conception, but a special kind of presence which Noë calls presence-as-access.

The idea of presence-as-access is best illustrated by how it accounts for a feature of visual phenomenology called presence-in-absence. Noë's favourite example features a tomato (2004; 2012). As with any opaque, three-dimensional object, you can never see every side of a tomato at once - you see one side, while another side remains out of view. However, Noë claims that when you look at a tomato, the rear side is nonetheless perceptually present 'as absent'. This is motivated by attention to visual phenomenology. Noë takes it as a primitive that the back of the tomato has perceptual presence even though it is not visible in the way that the front of the tomato is, a claim he justifies in part by noting that it does not seem as if you are merely looking at a two-dimensional cut-out of a tomato. The perceptual presence of the back of the tomato is enabled, he suggests, by your understanding of the movements that would bring the rear of the tomato into view.

The perceptual presence of objects and properties that are out of view is a pervasive feature of perceptual experience, and Noë emphasises the distinctive ability of presence-as-access to account for it. Another central case is the perceptual presence of

occluded parts of an object, for example a cat sitting behind a fence (Noë, 2002; 2004). Although parts of the cat are occluded by the fence posts, a true characterisation of visual phenomenology will observe that a visual experience of the cat is not an experience of a cat's parts, but of a whole cat. This is accounted for, as before, by the perceiver's understanding of the movements they or the cat would have to make to bring the occluded parts into view.

Presence-as-access is also intended to account more generally for the perceptual presence of a richly detailed visual environment. Although the visual system does not process a lot of environmental detail all at once, as evidenced by the experiments on change blindness and inattention blindness, Noë (2002) suggests that visual perceivers experience the environment's rich detail as present thanks to their implicit understanding of the movements, such as eye movements, that would bring relevant detail into view. Interestingly, he denies that perceivers ordinarily suffer from any delusion about the character of their access to environmental detail (cf., Dennett, 2001). He suggests that the surprise participants in Simons and Chabris' (1999) inattention blindness experiment reported on finding out that they failed to notice the gorilla walking through the scene is explained by the implicit expectation that their attention would be directed toward the gorilla, rather than the belief that visual experience presents rich detail all at once.

Noë (2012) claims that there is no firm dividing line between objects that are perceptually present and objects that are not. He claims that a perceiver may even have quasi-perceptual experience of an object in the very distant environment, such as a friend living in a distant city. If you know where that friend lives, your thoughts about him are characterised, in part, by an understanding of the fact he is far away and an understanding of the movements you could make to reach him, such as getting on a plane. For this reason, the thought takes on, in part, a quasi-perceptual character, although Noë does not claim this counts as a proper instance of perceptual presence. He contrasts this with the case of a friend with whom you have lost touch. Because you do not know where the friend lives, or whether he is alive or dead, the phenomenology of thoughts about this friend is quite different, a difference explained at least in part by the fact that thoughts in the latter case do not draw on an understanding of your sensorimotor relation to the friend.

The most striking feature of presence-in-absence (the phenomenology) taken together

with presence-as-access (the logical claim that accompanies it) is that all perceptual experience is supposed to belong to this type. This means presence-in-absence and -as-access incorporates experiences we might be tempted to think feature a simpler kind of presence, for instance the experience of the front side of the tomato. This is a counter-intuitive and, on a certain reading, highly problematic claim, so it is worth noting that Noë is explicit on this point. He says: “*all* perceptual presence is presence as absence [...] perceptual presence is availability; the world shows up for us as accessible” (Noë, 2012, pp. 95-96). Instead of imagining that the tomato’s facing side has presence simpliciter while the non-visible parts are presented as absent but accessible, Noë is explicit that the perceptual presence of the facing and the rear side both belong to the latter kind. This is supposed to be evident from introspectable features of visual phenomenology. In a description duplicated with slight modification from a passage originally found in his earlier book, he says:

[S]top and look again, for example, at the face of the tomato. You can’t comprehend the whole of it all at once in your visual consciousness. You focus on the color now, but in doing so, you fail to pay attention to the shape, or to the variations in brightness across the surface, but only at the price of ignoring the rest of it. You can no more achieve perceptual consciousness of the whole aspect of the tomato’s front side all at once than you can see the tomato from every side all at once [...]. In a way, for perception, everything is hidden. Nothing is given. (Noë, 2012, p. 19)

The observation that we cannot fixate on atomic features of the visual environment is compelling, but the suggestion that everything is hidden is more obscure and can be glossed in a number of different ways, as we will see.

In his earlier account, Noë (2004) gives a similar description of visual phenomenology and suggests it falls out of a particular understanding of the physical process that underlies perceptual consciousness. To this effect, he offers an analogy between visual perception and a computer accessing an online version of a newspaper via the World Wide Web. Rather than downloading the entire issue at once, the computer downloads the articles one at a time, accessing them only when called upon by the user. Downloading the issue in a piecemeal way has certain advantages. It limits the burden placed on the computer’s bandwidth and processing power, and it means that if an article is amended, the user can rely on seeing the most up-to-date version. Notice that this is similar to perception, understood in accordance

with the ‘outside memory’ thesis that perceivers use eye movements to access missing detail, when required, directly from the outside world. Importantly, it makes no practical difference to the person reading the newspaper that the computer downloads the articles one-by-one rather than all at once, since they have immediate access to the articles as they request them. As Noë puts it, the computer and its user therefore have ‘virtual’ access to the entire newspaper. This is comparable to conscious vision, in which he alleges not only that perceivers access detail in the way suggested by the outside memory thesis, but also that they have a phenomenological sense of that detail’s presence by virtue of knowing how to access it when required. Foreshadowing the claim that *all* perceptual presence is presence-as-access, Noë claims that conscious vision is in part disanalogous to the computer case, because rather than factoring into an offline and online component, perceptual presence is “virtual *all the way in*” (2004, p. 134).

The computer analogy may invite a problematic understanding of the claim that in perception, everything is hidden. While the computer ‘knows’ how to access the New York Times, it does not anticipate, in advance, what the contents of the articles will be. Noë’s descriptions of presence-as-access frequently imply that conscious accessibility works in a similar way. On a certain reading, this suggests that subjects never consciously perceive the environment, and that all they are conscious of is a disposition to make exploratory bodily movements. The claim that perceptual consciousness consists of no more than this is suggestive of an outright denial of perceptual consciousness. Noë (2004) explicitly disavows the suggestion that denying perceptual consciousness is a consequence of the account. If it was a consequence, it would be difficult to accept, since it is not clear that there is good reason to think that conscious perception encompasses knowing how to access missing detail but does not encompass the detail itself. Moreover, if the theory deflates qualia so much that phenomenal consciousness is no longer supposed to exist at all, the sensorimotor account loses much of its explanatory advantage when compared with accounts of neural-correlates of consciousness, which themselves invite the suggestion that there is nothing more to phenomenal consciousness than being in a particular neural state. Although it is a conceptual possibility that we are radically deceived about the existence of perceptual consciousness, the claim is not a good fit with the sensorimotor theory.

There are, however, more promising ways to make sense of the claim that everything is hidden. At times, in particular in his description of P-property experience, Noë (2004) implies that understanding which movements will bring an object (or part of an object) into view is more than merely knowing how to access missing detail. It is equivalent to knowing what that object is actually like. As we saw earlier, knowing what movements to make to face an object involves knowing the object's relative location, and having the ability to reach out and grasp it, for instance, implies knowing something about its perspectival shape. With this in mind, presence-as-access could entail, in a disanalogy with the online newspaper case, that perceivers have somewhat detailed prior expectations about what the results of their bodily engagements with the environment will be. These expectations may be enough to explain the presence-in-absence (rather than mere absence) of determinate features of the visual environment. Moreover, there could be nothing more to perceptual experience than these expectations.

The claim is plausible when we consider the analogy with touch, favoured especially by O'Regan (1992; 2011). You close your eyes and pick up an object, for instance a bottle. Your feeling of holding a whole bottle, despite only being in sensory contact with a small part of it, is explained by expectations about the results of possible movements. At the same time, you have a less specific expectation that reaching out to touch other parts of the environment will reveal the presence of an assortment of further objects. The presence of the bottle, and the more general presence of a detailed environment, are genuine examples of conscious presence and are each accounted for by sensorimotor mastery, even though they are slightly different phenomena. It might seem that there is also a narrower kind of presence which does not consist of sensorimotor expectations but merely of a conscious sensation caused by parts of the environment you are in sensory contact with, for example, the feeling of pressure on your fingertips where they are in contact with the bottle. There is certainly a difference between your awareness of the parts of the bottle that are in contact with your skin and those that are not - I will examine Noë's explanation of this in a moment. What is in question, however, is that the experience of pressure on your fingertips is anything other than the expectation that given certain movements the input to your sense organs (O&N, O'Regan, 2011) or your alignment to the object (Noë, 2004; 2011) would vary in systematic ways. The sensorimotor

account contends that once we discount these expectations, there is no raw feel - as Noë would put it, no qualitative residue - remaining.

One might be tempted to think that perceptual consciousness divides into a sensational component, which might be generated by inputs from the eyes and other sense organs, and a separate intentional component that gives sensations their world-directed content. Noë denies that sensations account for phenomenal character. He says: “We can’t explain the qualitative features of experiential states (veridical or nonveridical) by supposing that they depend on the instantiation of sensations in consciousness” (2012, p. 63), and elsewhere: “Where there is an object, but no [sensorimotor] understanding, there is nothing that even rises to the level of being misleadingly like perceptual consciousness; there is only, in effect, blindness” (2012, p. 25). This might be taken to allow that perceptual experience does factor into distinct components, namely sensation and sensorimotor mastery, but that sensorimotor mastery is a necessary condition for perceptual experience.

This would be one way of explaining why your experience of the parts of the bottle that are in contact with your skin seems different from your experience of the parts of the bottle that are not. However, given that sensations cannot explain the phenomenology or intentionality of perceptual experience without sensorimotor mastery, it is hard to see why they should play a role once sensorimotor mastery has entered the story. Thus we should take it that the feel of the parts of the bottle in contact with your fingertips is explained by your grasp of how your relation to the bottle would change in line with movement, and does not depend on anything distinct from Sensorimotor Understanding. The present story makes a good fit with the claim that Sensorimotor Understanding is realised by the deployment, subpersonally, of representations of SMCs (a claim found, for instance, in a proposal by Seth, 2014, more on which in chapter 4). It is also compatible, in principle, with the suggestion that perceptual consciousness is enabled by nothing more than a non-representational dynamical system.

There is a slightly different way to understand the claim that all presence is presence-in-absence, which suggests that conscious perception does not just depend on possessing practical knowledge about the movements you could make to bring missing detail into view, but also, in an ineliminable way, on the way that the knowledge is exercised in the

course of a perceiver's temporally-extended engagement with the environment. Noë does not unequivocally endorse this suggestion, but hints at this view by observing that perception is a temporally-extended process (2004; 2012). As Clark (2009) comments, the way the process unfolds over time may be essential to the character of perceptual awareness, meaning that there may be nothing about a perceiver's physical state considered *at a fixed moment in time* that would yield perceptual awareness with a determinate content or phenomenal character. If this were true, it would be possible to say, with reference to the perceiver's bodily dispositions, which aspects of the environment the perceiver, at a given moment, has the possibility of accessing. However, we would have to refer to the perceiver's temporally-extended engagements with the environment, which might incorporate the exercise of those dispositions, to fully account for the character of her perceptual experience. I make a suggestion along these lines in chapter 7.

I have one final point to discuss before concluding this section. Noë (2012) observes that there is a need to distinguish the kinds of perceptual presence enjoyed respectively by the front of the tomato, the back of the tomato, and objects in another room. He claims that perceptual presence admits of degrees, such that the front of the tomato has more presence than the back of the tomato, while objects in the next room are not, in a strict sense, perceptually present at all. To account for this, he appeals to quantifiable properties called movement-dependence and object-dependence, each of which are necessary for full-blown perceptual presence to occur. The former is the degree to which movements by the perceiver 'manifestly' change the perceiver's sensorimotor relation to a feature of the outside environment, while the latter is the degree to which movements by the object do the same.

Assume, for the time being, that a 'manifest' change in the sensorimotor relation is one that coincides with a change in retinal stimulation. I will return to this point in a moment. Noë observes that small movements by the eyes or the tomato are enough to manifestly change the sensorimotor relation between the perceiver and the front of the tomato, meaning this relation bears a high degree both of movement-dependence and object-dependence. A small rotation by the tomato is enough to manifestly change the perceiver's sensorimotor relation with the back of the tomato, but something greater than an eye movement is required by the perceiver to effect a manifest change in this case - for instance, she might have to stand

up and walk to the opposite side of the table. This means that the subject's sensorimotor relation to the back of the tomato bears a high degree of object-dependence and a more limited degree of movement-dependence. This accounts for the lesser perceptual presence of the back of the tomato. The perceiver's sensorimotor relations to objects in other rooms, or other cities (etc.), bear only negligible degrees of object-dependence and movement-dependence, since dramatic movements by the perceiver or object are required to effect changes in the retinal stimulation the subject receives from those objects. The degree of object- and movement- dependence in these cases is sufficiently low to entail that the objects are not perceptually present at all, although they may, as we have seen, bear a kind of quasi-perceptual presence.

It is apparent that Noë has to stipulate that the movements in question cause 'manifest' changes in the character of the sensorimotor relation, since a perceiver stands in a sensorimotor relation to every object in the world, each of which is equally prone to changing as the object or perceiver moves. However, he does not explain what it takes for a change in the relation to be manifest in the appropriate sense. Notice that in some cases, we might understand the way movements cause sensorimotor relations to change even when they are relations with unperceived objects. For example, you could be in Europe walking due north and reflecting on the fact that every step forward takes you an equal distance closer to the North Pole. In the event that the position of the North Pole suddenly changes, a compass needle will be enough to show you how your sensorimotor relation with the North Pole has changed. This means your relation with the North Pole satisfies the criteria for object-dependence and movement-dependence, somehow construed. This would not, of course, count as a case in which the North Pole is perceptually present. This is not intended as a counterexample to Noë's claims, but simply to underline that an explanation must be given of what it takes for a change in a sensorimotor relation to be manifest in a distinctively perceptual sense.

This is important, because the invocation of a perceptually manifest change to the sensorimotor relation invites a misunderstanding of the sensorimotor theory. Specifically, it invites the faulty view that the changes to the sensorimotor relation become manifest first, and that this subsequently causes the perceiver to revise their understanding of the

sensorimotor relation. One understanding might be that objects cause us to have conscious sensations, and that the sensations come to be intentionally-directed when we understand how they change in line with movement. As I have noted, Noë rejects this understanding of the sensorimotor account: rightly so, as it would undermine the insight that phenomenal character can be accounted for naturalistically by nothing more than a skill-driven sensorimotor relation to the outside world. An alternative understanding of the sensorimotor theory might say that the world is, in fact, represented pictorially in perceptual consciousness, and that changes to this pictorial representation enable you to revise your understanding of the sensorimotor relations you stand in. This is no good, because sensorimotor understanding is required for perception to be intentionally-directed in the first place, and we could not therefore appeal to the world-directed content of pictorial representations to account for sensorimotor understanding.

The proper understanding of ‘manifest’, I suggest, instead specifies that the perceiver does not consciously access whatever change in input causes the change in the sensorimotor relation to be manifest. On this basis, we should take it that the relevant input is a subdoxastic sense input. The upshot, given that movement-dependence and object-dependence are necessary conditions for the perceptual presence of P-properties and objective properties, is that a necessary condition for perceptual presence, at the personal level, is that subdoxastic sense inputs originating from the objects and properties consciously perceived must currently be prone to exhibiting patterns of movement-related change.

Notice that the claim that perceptual experience does not factor into an occurrent and a merely potential part is, on the present understanding, compatible with the outside memory thesis (O’Regan, 1992), which claims that visual processing involves an occurrent part, namely the retinal image, and a potential part, namely expectations - represented, or physically-constituted in a non-representational way - about the consequences of possible movements. The current personal level account can be taken to imply that this distinction between the occurrent and the merely potential does not manifest itself in a straightforward way to the conscious agent. This makes sense when we consider that appeals to the character of retinal stimulation are not sufficient to explain the phenomenal character or intentionality of perceptual consciousness. In light of this limitation, Sensorimotor Understanding is all

there is to go on. Noë's notion of presence-as-access is in principle compatible, unlike the outside memory thesis, with the existence subpersonally of filling-in and pictorial representation, since it does not need to claim that there is a one-to-one mapping between perceptual experience and the physical processes that enable it. Indeed, Pessoa et al. (1998) and Noë and Thompson (2004) actively deny that there is such a mapping between subpersonal content and the contents of visual experience. Notice, all the same, that presence-as-access does not do anything to invite the view that filling-in or pictorial representation take place subpersonally, and makes a very good fit with the outside memory thesis.

3.5 - Perception and Dynamical Entanglement

According to one current within sensorimotor theory, associated in particular with Hurley (1998), perception should not be understood subpersonally as a matter of information processing, but as a dynamic process comprising a messy entanglement between different parts of the brain, non-neural body and outside environment. Hurley contrasts her view of perception and action with an orthodox conception she calls the 'classical sandwich', which claims that the mind decomposes into discrete subpersonal systems responsible, respectively, for perception, thought, and action. Perception and action are the respective top and bottom layers of the sandwich, and thought makes up the filling. Thus perceptually-guided action, on the classical sandwich conception, takes the form of a 'sense-represent-plan-move' cycle (Wheeler, 2005). The visual system uses retinal inputs to construct a representation of the outside world, in a manner broadly in line with the approach offered by Marr (1982). This representation serves as an input to an informationally encapsulated and functionally distinct central cognitive system. The central system chooses a course of action and sends instructions to a separate motor action system, which is responsible for generating the appropriate motor movements. Fodor's (1983) account is a good example of this kind of framework.

One's preferred conception of the causal processes involved in action and perception can be laid out diagrammatically with sense inputs at the top and motor outputs at the bottom. Hurley suggests that the classical conception endorses 'vertical' modularity, since the

modular layers responsible respectively for perception, cognition and action are stacked vertically on top of one another, with perception at the sense input end and action at the motor output end. Hurley agrees that the mind decomposes into a number of informationally encapsulated systems geared toward differing tasks, for example different kinds of perceptually-guided action. The layers in her conception, however, are stacked side-by-side, meaning the boundaries between sense inputs, motor outputs and central cognition are non-existent or at least considerably more porous than conceived by the classical sandwich view. Thus one distinctive feature of perception in Hurley's account is that it depends constitutively on fluidly integrated causal processes running between sense inputs and motor outputs.¹¹

Hurley suggests that the causal flow runs from sense input to motor output and, importantly, loops back to sense input again, as body movements alter the character of the sensory stimulation received from the environment. In addition, she claims that the causal flow runs in the opposite direction too, as efferent copy from the motor outputs makes a causal impact upstream. She states that these points distinguish her account from behaviourism, which conceives of cognition as depending solely on unidirectional relations between sense inputs and motor outputs. Thus Hurley claims perception and action depend constitutively on messy dynamical processes, carried out without internal representations¹², and involving a complex entanglement between different parts of the brain, extra-neural body and environment. In addition to claiming that the substrate of perceptual consciousness is wider than sometimes thought, incorporating motor regions of the brain and perhaps the non-neural body and environment, Hurley claims that perceptual experience depends constitutively on the way the interactive process unfolds over time. Thus she rejects temporal

¹¹ Chemero and Cordeiro (2000) object to Hurley's continued use of the words 'input' and 'output' on the ground that on the dynamical conception, which they and Hurley endorse, there is no principled reason for conceiving of them as inputs and outputs. Still, we can describe them as inputs and outputs for heuristic purposes while acknowledging that this may not reflect a genuine functional distinction.

¹² As with 'inputs' and 'outputs', Hurley retains an apparently orthodox commitment by endorsing the vehicle/content distinction, implying that the entire processing loop could be considered one big representation. This is nonetheless quite different from the wholly internal representations endorsed by the classical sandwich conception.

atomism, the claim that the processing can for explanatory purposes be broken down into temporally discrete chunks.

With the points in mind, perception (and action) are held by Hurley's account to be best interpreted as dynamical systems. As such, they have the potential to be modelled using dynamical systems theory, which uses difference or differential equations to describe the ways a mechanistic system evolves over time. Hurley's account is just one example of a broader 'dynamicist' approach to cognitive science, associated in particular with van Gelder (Port and van Gelder, 1995; van Gelder, 1998) and Thelen and Smith (1994) and encompassing work in biological enactivism (e.g., Thompson and Varela, 2001) and ecological perception (Chemero, 2009). All endorse the suggestion that cognitive activities can be best understood, and in particular formally modelled, as temporally-evolving, dynamical processes, usually without reference to representation.

Aside from Hurley's account, canonical work on the sensorimotor theory is by and large non-committal about the dynamical approach. O'Regan (2011) makes no mention of dynamical systems, appearing, in spite of his account's emphasis on interactive engagements with the outside environment, to prefer an information processing stance. Noë (2004) only claims that perception *might* depend constitutively on dynamic input-output-input loops. The dynamical view is given a fuller endorsement, however, by other work on the sensorimotor theory, including Hurley and Noë (2003), Gangopadhyay and Kiverstein (2009) and Burhmann, Di Paolo and Barandiaran (2013).

There are numerous available motivations for endorsing the claim that perception depends constitutively on an embodied dynamical process. One such motivation is evidence that the phenomenal character of perceptual consciousness is partly determined by motor output activity. Standard representationalist models do not expect perceptual experience to be non-instrumentally dependent on activity by the brain's motor regions. If perceptual experience depends constitutively on motor activity, it plausibly depends on entire loops of interaction between brain, body and outside environment. Hurley (1998) appeals to an experiment by Kohler (1962) featuring goggles in which the left half of each lens was tinted yellow and the right half tinted blue. After wearing the goggles for several weeks, the subjects adapted to the coloured lenses, meaning that the visual scene came to appear normal and

untinted. After removing the glasses, the subjects reported that moving their eyes to the left caused objects to appear with a yellow tint, while moving their eyes to the right produced a blue tint. Hurley claims that this supports the view that motor output states partly drive visual experience and determine its phenomenal character.

The dynamical approach can also be motivated by the idea of the world as an outside memory. Although O'Regan does not pursue the outside memory idea to this conclusion, the idea that perception is best characterised as a dynamical system is a natural extension of the claim that perception depends on embodied interactions and not solely or primarily on internal information processing. Thus the arguments provided by O'Regan and others in favour of the outside memory view, for example the empirical work on change blindness and inattention blindness, give some evidential support for a dynamical systems version of the sensorimotor theory.

Elsewhere, the dynamicist approach is strongly associated with accounts that claim internal representations do not exist at all or cannot do the explanatory work required to bridge the explanatory gap between consciousness and the physical. As Chemero (2009) comments, dynamical systems models can frequently be given a representationalist gloss, but such a gloss is not necessarily required for the models to do the explanatory work they do. This contrasts with orthodox information processing accounts, which typically make representation an essential posit. Hutto and Myin's account (2012), which claims on conceptual grounds that subpersonal representational content probably does not exist, thus provides support for the dynamical systems perspective. Biological enactivists (Varela, Thompson and Rosch, 1991; Thompson, 2007) have a further reason to endorse dynamical systems, namely the view that cognition is continuous with the self-organising dynamical systems they hold are constitutive of life.

Chemero (2009) claims that the merits of a nonrepresentational dynamical systems approach must be judged by its explanatory success in making sense of empirical data rather than by metaphysical considerations about content of the kind endorsed by Hutto and Myin. However, Silberstein and Chemero (2012) offer a metaphysical-motivation of sorts for adopting a dynamical systems perspective, claiming that it helps motivate a direct realist account of phenomenal qualities, which in turn means it promises to do a better job than

competing accounts of explaining phenomenal qualities physicalistically. I will return to this point in the next section.

3.6 - The Extensive Conscious Mind

The sensorimotor account, in some variants, advocates one of the more extreme departures offered recently from the mainstream view that the mind is located purely 'in the head' and does not need to be explained by appeal to the body or outside environment. The orthodox position, to which the sensorimotor theory sets itself in opposition, is sometimes traced back to Descartes, who claimed that the mind was an immaterial entity distinct from the body, with which it interacted via a prescribed channel at the pineal gland. Since the mind, on the Cartesian view, is not even physical, it is in this sense more disengaged from the body and physical environment than it ever is in contemporary cognitive science. However, orthodox physicalist accounts of mind continue to subscribe to what Wheeler (2005) describes as Cartesian psychology, an approach which implies, amongst other things, an analogous separation between the neurally-realised mind and the extra-neural body. Descartes was, in fairness, careful to claim that the mind is intimately joined with the body and not detached, as he put it, in the way a pilot is from their ship (1647/1985b). Wheeler, attempting to do justice to the position, suggests this claim can be reconciled with Descartes' dualism if we regard it as a principle of Cartesian psychology that the mind makes use of informational content specifying the body's state, but that the way it uses this information to generate intelligent action can be fully characterised without reference to the body. This principle, drawn from a nuanced version of Cartesian psychology, is explicitly endorsed by a number of recent, largely orthodox accounts of cognition, which hold that the mind supervenes purely on the brain but frequently trades in representations of bodily states (e.g., Goldman and de Vignemont, 2009; Rupert, 2011).

Drawing a less kind parallel between Descartes and contemporary accounts, Dennett (1991) issued a well-known warning against being tacitly committed to Cartesian materialism, the view that the mind is unitary and its physical substrate confined to a circumscribed area of the brain - an outlook that must be avoided on pain of committing the homunculus fallacy.

Present day theories often maintain an orthodox position while appearing to avoid this mistake. Rupert (2011) is again a good example, since it is the central thesis of his ‘massively representational mind’ that representations are manifold throughout the brain, although he maintains that cognition is overwhelmingly brain-bound. Noë (2004) argues that brain-bound accounts cannot avoid the Cartesian mistake altogether, since they continue to endorse a conception of the brain itself as a homunculus subsisting within the wider confines of the person as a whole. This need not be viewed as a mistake, so long as you are happy to endorse the view that in an important sense, you are your brain. This is nonetheless strikingly similar to Descartes’ claim, with reference to his immaterial mind, that “I am in the strict sense only a thing that thinks, that is, I am a mind or intelligence or intellect or reason” (1647/1985b, p. 18). The moral is that although orthodox theories are often deliberate in avoiding the extremes of Cartesianism in a crude rendition, they retain Cartesian commitments.

One point of departure from the Cartesian orthodoxy occurs with the thesis known as vehicle externalism. To explain the generic position, it will help to compare it with content externalism (Putnam, 1975; Burge, 1979). This longer-established position holds that mental content is not fixed solely by things in the head. In Putnam’s classic thought experiment, Oscar, who lives on Earth, has an exact physical duplicate, Twin Oscar - who lives on Twin Earth. Twin Earth is identical to Earth in every respect, except the clear liquid which fills rivers, lakes, and has all the other surface features we associate with water, has the chemical composition XYZ rather than H²O as it does on Earth. Oscar and Twin Oscar live in medieval times and are ignorant of the microstructural properties of water. But when Oscar and Twin Oscar each utter the word ‘water’, the differing properties of the natural kinds they are in causal contact with entails that they express different thoughts.¹³ The moral is that identical vehicles can bear non-identical contents, because content is fixed in part by properties of the outside environment. Vehicle externalism bears comparison with content externalism, although it is a quite different position. Importantly, content externalism is not

¹³ When Putnam (1975) formulated the view, he intended it to apply to linguistic rather than mental content, but the thesis came later to be widely applied by Burge (1982) and others to the thoughts underlying the utterance and not just the words themselves

generally thought to threaten the view that the vehicle of thought and experience is wholly neural. Vehicle externalism, by contrast, contends that the physical substrate of contentful mental states is sometimes outside the head.

Vehicle externalism, generically, decomposes into a number of quite distinct and, indeed, frequently incompatible strands. One such position is the extended mind thesis (Clark and Chalmers, 1998; Clark, 2008). Proponents of the extended mind frequently endorse functionalism (Clark, 2008; Wheeler, 2010), and hold that although cognition frequently involves the construction and deployment of representations, typically in the brain, it sometimes makes use of representational vehicles outside the skin. The most famous example, from Clark and Chalmers (1998), uses a 'parity'-based heuristic to lend support to the extended mind. Two protagonists, Otto and Inga, go about their daily business. But Otto, unlike Inga, suffers from mild dementia, meaning that unaided he suffers from severely impaired short term memory. To compensate, he makes use of a notebook which he keeps with him at all times, frequently noting down relevant facts and occurrences from his daily life, and calling upon the notes automatically and non-critically as needed. By interacting with the pen and notebook, Otto is able in a relatively smooth way to replicate the mental function for which Inga relies on her neurally-encoded short term memories. Clark and Chalmers suggest that because representations in the notebook, in Otto's case, function in an appropriately similar way to the relevant neural representations for Inga - hence the parity - the pen and notebook can be regarded as forming part of an extended cognitive system.

The extended mind thesis is one way of breaking from the tradition which says that the physical substrate of cognition is purely internal. However, the thesis is conservative in several key respects. It maintains a commitment to internal representation, and indeed uses the idea that cognition frequently depends on internal representation to motivate the parity-based claim that Otto's cognition extends, given that he uses representations in the notebook in the way other cognisers use internal representations. Moreover, it is committed only to the claim that unconscious mental processes sometimes extend beyond the skin, and proponents of the extended mind are sometimes explicit in denying that the substrate of consciousness extends (e.g., Clark, 2009).

The extended mind contrasts with a quite different vehicle externalist approach

associated with sensorimotor theory and other dynamical-systems oriented accounts of perception, which claim that the substrate of perceptual consciousness incorporates the extra-neural body and environment in addition to the brain (Hurley, 1998; Noë, 2004; Rowlands, 2010; Silberstein and Chemero, 2012). This approach, which I focus on for the remainder of the section, is sometimes called the ‘extended conscious mind’, a label intended to capture its resemblance to the extended mind. In fact, this label may not have the appropriate resonance, since it suggests a continuity with the extended mind thesis that does not reflect its quite different motivation. Vehicle externalism, strictly speaking, is a problematic label too, since some of the views falling under the present approach reject the claim that cognition, and specifically perception, involves content at all, and hence that there is a vehicle/content distinction to be made in this case. After Rowlands (2010), we could recast the distinction between vehicles and contents as a distinction between vehicles of consciousness and the conscious states that supervene on them.

A better label is perhaps suggested by Hutto and Myin (2012), who describe their approach to cognition as the *extensive mind* to underline their view that because there is no content there is, in the first place, no theoretically relevant inner-outer boundary between the brain and its surroundings from which cognition might extend. Hutto and Myin are sympathetic but not fully signed-up to the claim that perceptual consciousness supervenes on extra-neural activity. To make sure it has all the correct resonances, let us call the present view the *extensive conscious mind* (ECM). The core principle underlying ECM is the claim that perception is best explained by appeal to dynamical systems rather than computation. Although this claim is motivated in varying ways, and in turn supposed to justify ECM on the basis of varying further premises, adherents of ECM consistently endorse the principle that perception is best understood as a dynamical system, in principle describable using differential equations, incorporating features of the brain, body and outside environment.

A large part of ECM’s appeal is its distinctive way of accounting naturalistically for consciousness and phenomenal qualities. Because the objects of perception are themselves constituents of perception, phenomenal qualities can be accounted for non-mysteriously by appeal to physical properties the objects bear in relation to the perceiver’s body - for example, SMCs. One place ECM is endorsed is in Silberstein and Chemero (2012, henceforth S&C),

who advocate a version of the ecological approach to perception. The approach differs from sensorimotor theory by claiming that the features we are responsive to in perception are not SMCs but affordances. Like SMCs, affordances are relations between a perceiver and their environment, and can be straightforwardly specified from a third-person perspective. Unlike SMCs, they describe possibilities for bodily movement rather than sensory consequences of possible movement. Notwithstanding this difference, S&C's arguments in favour of ECM are compatible with the sensorimotor theory.

S&C emphasise that ECM can dissolve the hard problem of consciousness, by which they mean the problem of accounting naturalistically for specific phenomenal qualities *and* the more general question of why there is consciousness at all. They take inspiration from a point made by Ryle (1979), who identified a dilemma faced by theories of discursive thought. Ryle claimed the theories were prone to pivot between two unsatisfactory positions, either casting thought as 'nothing but' inner speech, or casting it as 'something else as well'. He illustrated the problem by analogy with a coin. It would be wrong to claim that a coin is 'nothing but' a metal disc, but wrong to claim it consists of a metal disc plus some other entity as well. S&C argue that the problem facing theories of consciousness takes a similar shape. Proponents of orthodox approaches to cognition sometimes claim that phenomenal consciousness is 'nothing but' a certain kind of neural activity - an unsatisfying position which also comes under target in O&N, as I discussed in section 3.2.2. Alternatively, they accede to a dualist impulse and claim, like Chalmers (1996), that consciousness is a non-physical 'something else as well'.

S&C blame computationalists for the apparent dilemma, offering the following diagnosis: Computationalism, which as S&C take it entails the representational theory of mind, states that computation is sufficient for cognition. Computationalists remain loath, however, to accept that metal and plastic artifacts could be conscious merely by virtue of performing appropriate computations. They are therefore forced to separate consciousness from cognition, and for this reason account for consciousness separately, either deflating it to brain state or inflating it to something non-physical. To avoid the dilemma, S&C reject the representational theory of mind and instead propose that the properties to which we are intentionally-directed in perceptual consciousness are themselves constituents of perception.

The fact that the properties we perceive, affordances, are specified with egocentric reference to the perceiver is supposed to account for perception's intrinsic subjectivity, and therefore consciousness in general.

Although S&C's solution is appealing, it is worth taking note of a limitation they share with O&N. It is true that the representational theory of mind is not necessarily committed to the claim that consciousness, like cognition, is best explained by appeal to representation. However, there are also representationalist proposals which take seriously the claim that contentful states, meeting certain other functional criteria, are sufficient to account for phenomenal consciousness (e.g., Lycan, 1987; Dretske, 1995; Tye, 2000). Accounts in this category accord with S&C's demand that consciousness is neither 'nothing more' than brain state nor 'something else as well', and like S&C they attempt to avoid the dilemma by identifying consciousness with cognition. However, they do not use ECM to make this work. ECM has one key advantage compared to representationalism about phenomenal qualities, namely that it avoids much of the latter's explanatory burden. ECM does not have to offer (or hold out for) a naturalistic account of content, a project which faces serious conceptual obstacles (Hutto and Myin, 2012; see section 6.4). Notice that reducing phenomenal character to content only helps naturalise phenomenal character if the content can itself be naturalised. Moreover, the success of representationalist accounts of phenomenal character is in part dependent on the sort of content involved. For instance, externalistically individuated content faces a distinctive challenge in light of Twin Earth-style cases, as we will see shortly.

Hurley (1998), offering a different argument for ECM, endorses the vehicle/content distinction, arguing that the vehicle of conscious content extends beyond the brain. She claims that vehicles of content play an essential theoretical role in explaining the relation between the mental and the physical. Drawing a distinction between type-explanations and token-explanations, she claims the former explain why there is mentality at all. For example, in line with Millikan's (1984) biosemantics, facts about natural selection and the ancestral environment might explain how certain types of physical state bear content. To explain a subject's present mental state, we must also appeal to presently occurring states and processes. We are now offering a token explanation, because we are specifying which physical states and processes token the relevant type. Thus Hurley claims that when we talk about a vehicle of

content, we are attempting to offer a token-explanation.

If one wanted to deny that perceptual experience is contentful or depends on content-bearing vehicles, the vehicle/content distinction could be recast as a relatively harmless vehicle/experience distinction. This avoids the need to stipulate that experience involves truth or correctness conditions. However, maintaining the reformulated distinction is only valuable if phenomenal states are multiply realisable, meaning that non-identical physical processes can yield phenomenologically indistinguishable experiences. S&C maintain that the dynamical systems which constitute perception are ‘soft assembled’ from varying physical components, sometimes wholly internal and at other times incorporating varying parts of the outside environment. This implies that perception, in general, is in one sense multiply realisable. However, the point here is that the vehicle of perceptual experience in general can vary, bringing with it attendant differences in perceptual phenomenology, and not that an identical experience could supervene on a range of different vehicles. The sensorimotor theory, on the gloss I argue for later in the thesis, allows that perceptual experiences may in a weak sense be multiply realisable, given that phenomenal character could stay the same across perceivers with limited differences in body morphology. However, the multiple realisability that applies in this case is too constrained to make a vehicle/experience distinction particularly useful, especially when contrasted with the functionalist thesis that phenomenology could stay the same despite radical differences in embodiment.

Hurley’s arguments in favour of vehicle externalism are equally applicable to other versions of ECM. Conceding that neural processes are likely to be sufficient for consciousness, she challenges the assumption that suitable attendant outer states are not themselves necessary for the neural processes in question to occur. The ‘duplication assumption’, as she calls it, is a standard background assumption underlying the twin test for content externalism. Recall the Twin Earth thought experiment from Putnam (1975). Block (1990b) offers a similar thought experiment intended to test representationalist accounts of phenomenal character (see chapter 5). Suppose there is a planet called Inverted Earth, which is like Earth except the colours of objects are all inverted. A subject travels from Earth to Twin Earth and puts on colour-inverting goggles. As a result, the subject’s phenomenal states remain as they were on Earth, even though, according to content externalist accounts, the content of the subject’s

perceptual states must differ. According to Block, this demonstrates that content externalist variants of representationalism about phenomenal qualities fail. His own view is that we should appeal to some non-representational property of the brain to explain phenomenal character. Hurley's argument undercuts this debate by casting doubt on the claim that twin tests are possible. If successful, it challenges Block's view that we must only appeal to the brain, since it sheds doubt on the claim that phenomenal states can stay the same while the environment radically changes. It challenges the content externalist proposal, since it denies that the vehicle can stay the same while the content changes. Notice that it thereby offers a response to representationalism about phenomenal qualities absent from S&C's more recent piece.

Hurley observed that the duplication assumption becomes problematic when we apply it to realistic rather than toy cases. In *Consciousness in Action*, she offers a series of Twin Earth-style scenarios which challenge the assumption in progressively demanding ways. A first concerns a case of colour inversion in which all objects that are red on Earth are green on Twin Earth but *not* vice versa. Because the transformation between Earth and Twin Earth is not information preserving, the subject on Twin Earth is unable to make the same range of discriminations as the subject on Earth, therefore their neural states are not duplicated and the duplication assumption fails. This shows that even for cases where the differences between the Earths are limited to colour and not shape, the duplication assumption only holds for a heavily restricted set of cases.

A further scenario features a planet called Mirror Earth in which everything is reversed from left to right, so that everything is flipped on a vertical axis. The Mirror Earth subject wears left-right inverting goggles, meaning their retinal input, right now, is identical to the state it would be in on Earth. However, when the twin moves their right hand, the visual feedback they receive duplicates the visual feedback they would receive from moving their left hand, not their right, on Earth, and therefore duplication fails. To correct for this, we could imagine that a device is attached to the twin's peripheral nervous system, intervening in their motor outputs and proprioceptive inputs, such that the same instruction causes the subject to move their left hand, and hence receive the necessary visual feedback, while receiving the proprioceptive input they previously received after moving their right hand. This saves the

duplication assumption, for the time being, although it depends on confining duplication to the central rather than peripheral nervous system.

Lastly, there is El Greco Earth, in which everything is vertically elongated compared to Earth, such that the counterparts of spheres are egg-shaped. This leads to a challenging asymmetry between the respective Earths. When an Earth subject nudges a sphere, its orientation stays the same - perhaps it rolls a little to the left or right. The same action on El Greco Earth causes the egg-shaped object to fall over. Because of the asymmetry between the effects of motor movements on visual input in subjects on the two worlds, duplicating the Earth subject's internal states on El Greco Earth would require a vastly complicated device to correct the visual inputs in different situations. On a similar theme, Hurley remarks that real life Virtual Reality simulations depend on heavily restricting the range of possible movements afforded to the perceiver. The more freedom of movement the subject has, the greater their ability to disambiguate visual inputs by moving around, which means, claimed Hurley, that creating a fully realistic virtual reality simulation pushes toward the limit of what is, even in principle, physically possible. Noë (2004) makes a similar point, asserting that a completely realistic virtual reality simulation cannot be much less complex than the world itself.

The upshot, for Hurley, is that once it is established that the duplicationist must go to byzantine lengths, at the bounds of physical possibility, to preserve duplication with regard to a restricted area of the brain, the vehicle internalist thesis comes under threat. For one thing, there is no guarantee that the same area of the brain will be duplicated in each Twin Earth scenario. For another, even if some part of the brain can be consistently duplicated, without some strong motivation, the insistence that this region of the brain alone is the vehicle of consciousness appears arbitrary rather than principled. An important feature of this discussion is that it sets out to challenge duplication for a range of realistic cases. This does not mean denying that duplication is possible, even straightforward, in *some* cases. Establishing it in some cases is insufficient to show that perception in general supervenes only on neural activity. Hurley observes that in cases that are restricted enough to make duplication straightforward, for example because we imagine the subject's sensorimotor capacities are highly constrained, it is doubtful whether the subject enjoys perceptual experience in the first place.

She also pre-empts a number of more sophisticated objections that might be offered by a vehicle internalist. One such objection is that the laws of physics might vary between Earth and Twin Earth, such that complicated devices for altering the nervous system are not required to achieve duplication. Hurley counters that the twins cannot be considered duplicates if the laws of physics vary between the cases. Another objection is that a mad scientist could duplicate a subject's neural states in a different environment by intervening to alter them frame-by-frame. Hurley rejects this, because she rejects temporal atomism, the view that perception depends on a succession of moment by moment states rather than a smoothly continuous process.

The success of Hurley's argument, it should be noted, depends in part on endorsing a dynamical systems approach to perception. Indeed, Clark (2009) observes that it depends on more than just the claim that perception involves a temporally-extended, dynamical process. Such a process may merely be instrumental in causing the activation of a brain state which itself generates a perceptual experience in an instant, without requiring temporal extension. Alternatively, the brain state may need to persist through time, but it may not itself need to dynamically evolve through time in a way that would make the contribution of extra-neural states essential. Thus Clark suggests that Hurley's view must endorse a conception he calls Deep Dynamic Entanglement + Unique Temporal Signature, or 'DEUTS' (p. 980), which holds not only that perceptual consciousness is enabled by a temporally-extended, dynamic process, but also that the conscious experience is constitutively dependent on the precise way the neural as well as extra-neural states involved evolve over time.

If there are satisfactory independent grounds for thinking that perception depends on a specific functionally distinct, informationally encapsulated area of the brain, then it is no longer a prejudice to think that the same area of the brain can be duplicated in all the Twin Earth scenarios (notwithstanding the technical challenge) and that the area of the brain so duplicated is genuinely the area responsible for visual awareness. Of course, were the commitment to this view of perception motivated by a prejudice (recall the discussions of Cartesian materialism and analytic isomorphism) it would not be a suitable independent ground. Importantly, although a well-motivated version of vehicle internalism might support the duplication assumption, the duplicationist hypothesis cannot be used to motivate vehicle

internalism, as this would be getting things the wrong way round. Conversely, Hurley's argument shows that if perception depends constitutively on a dynamic, interactive process with a unique temporal signature, there is also a strong motivation for endorsing ECM.

Clark argues that if deep dynamic entanglement and a unique temporal signature are indeed essential to consciousness, this threatens a stalemate between ECM and its opponents. He argues that DEUTS cannot establish ECM, but it can expose the rejection of ECM as a prejudice. To show that ECM fails, the opponent has to identify a special property that the neural component of the processing possesses and that the extra-neural activity lacks, and show that this difference is a principled rather than arbitrary ground for thinking that perceptual consciousness supervenes solely on the neural activity. Clark claims that such a difference can be established in light of DEUTS' own appeal to a unique temporal signature. One aspect of the temporal signature that consciousness is likely to depend on is informational processing at very high speed. This, in turn, depends on the high-bandwidth available in the neural portion of the wider interaction between brain, body and environment. The interface between the brain and extra-neural body has a much lower bandwidth, suggesting it cannot do the work required to generate consciousness, and therefore serves as mere causal scaffolding (Clark, 2009). Importantly, Clark does not reject Hurley's insight that the neural processes involved in a particular conscious experience are likely to require that the subject is embedded in a specific type of environment, and hence her rejection of the possibility of duplication as illustrated by the Twin Earth-style cases.

I am not going to consider any further whether ECM finally goes through, as ECM is a mere nicety in the context of the sensorimotor theory as a whole. Equally important are the insights that lead up to ECM, including, on a dynamical version of the sensorimotor account, the claim that architecture of perception involves a dynamic, interactive process with a unique temporal signature, and the fact that the neural duplication is likely to fail. Silberstein and Chemero associate ECM's claim that the objects of perceptual awareness are constitutive features of perceptual processing with direct realism. However, it is not obvious that ECM is needed for direct realism to go through. Direct realism, including the relational realism endorsed by Noë (2012), could be motivated by the fact that perception is an interactive process in which duplication scenarios are impossible, without needing ECM too. Thus the

arguments that lead up to ECM are enough to get a dynamical version of the sensorimotor theory what it really needs.

4 - Mastering Sensorimotor Understanding

The sensorimotor theory's most crucial tenet is also the one least well-defined by the theory's flagship statements. Expanding on the work done in the previous chapter, I will in this chapter address in further detail the claim that perception depends on sensorimotor 'mastery'. Specifically, I will focus on the variety of sensorimotor mastery found in Noë's (2004) account, which I have officially labelled Sensorimotor Understanding, reflecting the terminology sometimes used by Noë. Many of the important things Noë says about Sensorimotor Understanding concern the 'presence' that arises when we exercise it (see section 3.4), and the nature of the properties it puts us in touch with, P-properties and objective properties (see section 3.2.5). In this chapter, I focus on Sensorimotor Understanding itself, with a view to addressing a lack of clarity about what it is, and how it fits into a scientific account of perception and perceptual consciousness.

Although the faculty featured in O'Regan (2011), which I label Sensorimotor Tracking, is characterised in somewhat different terms to Sensorimotor Understanding, the respective varieties of sensorimotor mastery are each loosely enough defined in the original work to leave open the possibility that they are one and the same faculty. Although my discussion will focus on Sensorimotor Understanding, the characterisation I end up with could also, therefore, describe Sensorimotor Tracking. I will continue this chapter by identifying what I regard as a key limitation with the concept presented by Noë, namely a lack of clarity about how to account for Sensorimotor Understanding physicalistically, particularly in the domain of cognitive science. In section 4.2, I address a specific ambiguity concerning the explanandum, and suggest two alternative ways of resolving it. In sections 4.3 and 4.4, I return to the scientific tractability issue by explaining the impact that the respective views of Sensorimotor Understanding should have on our understanding of sensorimotor theory as an account in cognitive science.

4.1 - 'Action in Perception' in Action in Perception

One of the most important differences between the early sensorimotor account (O&N) and

Noë's (2004) version is the leap from what we could think of as a 'bottom up' type of explanation to a 'top down' one. By bottom-up, I mean an explanation that starts by describing low-level physical states and processes that could be possessed by an unconscious machine, and expands to encompass consciousness and the personal level. By top-down, I mean an account that takes consciousness as its starting point and saves the question of how consciousness arises until later.

It will be useful if I explain Noë's motivation for this shift in emphasis. In the (2004) book, he notes that the early account focuses on how sense input signals, rather than conscious sensations, change in line with movement, observing that it therefore has no need to pre-suppose consciousness and is well-positioned to give a physicalistic account of how perceptual consciousness arises. The drawback with this strategy, Noë says, is that it comes at a cost to what he calls "phenomenological aptness" (Noë, 2004, p. 228). Expanding on this point, we might say something like this. The sense inputs featured in O&N are subdoxastic, i.e. inaccessible to the conscious agent. Even when the agent's mastery of the SMCs associated with those sense inputs is integrated with the appropriate capacities for thought and planning, the subject does not become aware of the sense inputs themselves, but of the relevant features of the outside environment. Since sensorimotor mastery, in the early account, is mastery of how sense inputs change with movement, there remains a question of which features of the outside world we consciously perceive, and how they show up for us.

The O&N account does answer these questions in part. The environmental features we perceive are those that cause the sense inputs whose patterns of movement-related change we, or our bodies and brains, have mastery over. The account also makes some observations about the phenomenal character of perceptual experience, and attempts in particular to resolve the problem of supposedly irreducible qualia by suggesting that there is nothing more to experiencing the redness of red, for example, than exercising one's mastery of the sensorimotor laws associated with redness for the purpose of thought and planning. However, this solution is in danger of failing to unseat the intuitions of someone who thinks that qualia are irreducible, because it invites them to ask why exercising mastery of a particular set of SMCs has one phenomenal feel and not another - this is another of the unresolved 'why' questions about the relation between consciousness and the physical that O&N make it their

project to avoid.

The problem arises because the O&N version of the sensorimotor theory cannot convincingly claim that perception *feels* like exercising sensorimotor mastery, given that the sense inputs which are the targets of the mastery remain subdoxastic. If perception does not feel like the exercise of sensorimotor mastery, the endorser of irreducible qualia will continue to hold that the relation between sensorimotor mastery and conscious feel is only contingent, meaning that the conscious feel is therefore nonidentical to the associated sensorimotor mastery. The sensorimotor theorist could respond that perceiving does feel like exercising sensorimotor mastery - necessarily, because perceiving just *is* exercising sensorimotor mastery - even though perceivers typically lack the ability to recognise that this is what perception feels like. However, if the sensorimotor theory concedes that conscious perception does not feel *recognisably* like the exercise of sensorimotor mastery from the subjective point of view of the perceiver, it finds itself deprived of powerful reason for thinking the sensorimotor theory is true, namely that the theory makes it impossible to conceive of a given act of perceiving feeling other than the way it feels. In any case, the early sensorimotor account simply does not go into as much detail about visual phenomenology as it could.

Like the original version of the theory, Noë (2004) claims that conscious perception is identical to the exercise of sensorimotor mastery. In a progression from the earlier account, however, he implies that introspection reveals that exercising sensorimotor mastery is recognisably what conscious perceiving feels like. This is useful because it provides a further reason for thinking that the act of consciously perceiving, with its distinctive qualitative feel, *just is* the exercise of a variety of sensorimotor mastery and nothing more. As we saw in the last chapter, the account claims that absent features like the back of a tomato are experientially present, in absence, because we know which movements would bring them into view. It also claims that objective properties, like the circularity of a plate, show up in visual experience as being mediated by the perspectival properties that presently obtain, which we implicitly understand will change in specific ways in line with possible movements.

In order to introduce these elements, Noë has to present sensorimotor mastery as a feature *of* perceptual consciousness rather than a faculty that enables perceptual consciousness. Because P-properties and objective properties are consciously accessible, it is natural to think

the faculty we use to access them is conscious, even though the understanding in question is implicit. Although Noë downplays the importance of the personal/subpersonal distinction, it is plausibly a requirement of Noë's account that the faculty we use to master P-properties and objective properties is itself a personal level faculty, since he explicitly intends the account to do justice to McDowell's (1994a) conceptualism about perceptual experience and Kant's famous maxim that "thoughts without content are empty; intuitions without concepts are blind" (CPR A51/B75), which in each case require that the *person* brings some capacities for thought to bear in experience. Noë does not say so, but it follows from this that Sensorimotor Understanding must also be amenable to a kind of self-reflexivity; not only do you consciously exercise Sensorimotor Understanding, but you have the capacity to be conscious *of* exercising it. This point notwithstanding, Sensorimotor Understanding is characterised by Noë most of the time as a conscious faculty possessed by the person rather than their brain or body, and Noë explicitly concedes that the account presupposes consciousness.

This reflects an important difference between Noë's sole-authored version of the sensorimotor theory and the others. As we have seen, the approaches to perception detailed in O&N and O'Regan (2011) focus, in particular, on processes that could be instantiated by an unconscious machine rather than a human perceiver. The accounts begin by describing a process of sensorimotor coupling sufficient for a certain kind of unconscious perceptual activity, and widen the explanation in a bottom-up fashion by introducing additional processes, the contributions of which they allege enable perceptual consciousness. Specifically, these are the integration of the lower-level process with thought and planning, and, in O'Regan's (2011) model, an additional faculty corresponding to the 'self'. Although the accounts do not give detailed proposals about how these higher faculties are physically realised, the authors intend us to suppose that the relevant faculties are things we could readily account for by appeal to appropriate physically-realizable states and processes. O&N at times suggest an information processing-style account is needed, while O'Regan (2011) explicitly endorses functionalism of a sort that means a robot could in principle possess the relevant functionally-defined faculties. Thus an important difference between these accounts and the version of the sensorimotor theory offered by Noë (2004; 2012) is that the latter pre-supposes consciousness and does not make any claim to explain it. The Sensorimotor Understanding

Noë appeals to is intended to explain something about perception and the character of perceptual consciousness, but is not thereby supposed to account for the existence of consciousness in the first place.

Neglecting to explain how consciousness in general arises is not an unreasonable limitation, as long as we think of the account as an account of perceptual consciousness rather than consciousness in general. More troubling is the fact that the account fails in great respect to commit itself to clearly defined theses about the physical processes implicated in Sensorimotor Understanding, something the sensorimotor theory could be expected to offer even in the absence of an account of how consciousness in general occurs. The omission would not be problematic if we took the account solely as an entry into philosophy of perception, which is not necessarily concerned with the physical states and processes that realise conscious perception. However, it becomes more of a problem when we consider that Noë continues to present the sensorimotor account as a rival to orthodox scientific accounts of vision. In response to Marr (1982), Noë says that: “Vision shouldn’t be thought of as a computation performed by the brain on inputs provided by the retina. What is vision? How should it be characterised computationally? This book suggests the outlines of an answer. *Vision is a mode of exploration of the environment drawing on implicit understanding of sensorimotor regularities*” (Noë, 2004, pp. 29-30). One of his points is that Marr presents the subpersonal processes that enable perception as if they were identical to the higher-level processes that perception is constituted by. This is not only meant to be a logical error, however. His objection is that the error causes Marr to get his account of the lower-level processes wrong. Indeed, Noë casts doubt on many of the principles needed to make a computational account stand up, claiming for example that perceptual consciousness depends in part of having the right kind of body (p. 25), a claim by which he intends something stronger than the mere suggestion that you need a body capable of implementing the appropriate computations.

We could simply take the sensorimotor theory as a rebuttal of orthodox scientific accounts like Marr’s without expecting it to provide an alternative account of the subpersonal processes involved in perception. It would be better, however, if the theory gave an account of the physical or physically-realised functional states implicated in Sensorimotor Understanding. Offering such an account involves answering some very abstract questions, for example

whether functionalism and representationalism are conceptually viable approaches for reducing the mental to the physical, and slightly more concrete questions, for example whether information theoretic or dynamical systems frameworks have the most theoretical utility for explaining what is going on in human brains and bodies. Even more concretely, it would be nice to know the precise mechanism involved in human Sensorimotor Understanding. For instance, if we concluded that the best way to physicalistically account for Sensorimotor Understanding is to suppose that it supervenes on the construction and deployment of neural representations, perhaps of some relatively novel kind, we would still need to know what kind of content they bear and what functional role they play.

Ward (2009), for example, has argued that the representations in question might specify possibilities for highly planned agent-level action. Elsewhere, Seth (2014) has proposed that aspects of the sensorimotor approach can be allied with predictive coding approaches to cognition, which I will explain below. In contrast with representationalist proposals like these, Hutto and Myin (2012) suggest that Sensorimotor Understanding must be explained by appeal to the agent's history of sensorimotor interactions rather than by representational content, although it remains to be properly explained how we should characterise the physical processes that underlie the exercise of sensorimotor mastery *right now* if we endorse this view. Burhmann, Di Paolo and Barandiaran (2013) have offered a way to model sensorimotor perception using a dynamic systems framework. The dynamical systems approach was also endorsed by Hurley (1998), whose account can be considered part of the sensorimotor approach, and Chemero (2009; see also, Silberstein and Chemero, 2012), who endorses a Gibsonian account of perception, but whose appeal to sensorimotor abilities offers a proposal that may usefully help do justice to the sensorimotor approach.

The founding statements of sensorimotor theory can be forgiven for not arriving with a prefabricated and fully-developed account of the physical processes involved in Sensorimotor Understanding, but if it is to convince as a serious rival to orthodox scientific accounts of vision, we need a direction of travel toward answering them. A good place to start is to have a clear working proposal about what the explanandum, Sensorimotor Understanding, actually is. In fact, even this so far is ambiguous.

4.2 - Which Direction of Fit?

In the sections that follow, I will further discuss the physical or functional processes involved in Sensorimotor Understanding, but in this section I will first address a significant ambiguity about what Sensorimotor Understanding is, namely the fact that it is sometimes pitched as Being Able to Act and sometimes as Knowing What Would Happen if You Did Act.

The former, Being Able to Act, could be understood in various ways. In a position approximating behaviourism, it could be taken to mean having a disposition or ability to act. Alternatively, in a position resembling Gibsonian ecological psychology (e.g., Chemero, 2009) or the literature on action-oriented representation (e.g., Clark, 1997; Wheeler, 2005; Ward, 2009; Ward, Roberts and Clark, 2011), it could be taken to mean knowing what actions and behaviours are open to you. Although Noë (2004) sometimes disowns these possibilities, there is much evidence that even his version of the sensorimotor theory is committed to a view of Sensorimotor Understanding as Being Able To Act. If Sensorimotor Understanding is like this, it has a *desire-like* direction of fit, because it is the state of the world, specifically the agent's behaviour (which for this purpose counts as a state of the world) that must fit with the agent's mental state, and not the other way round. Contrast this with Knowing What Would Happen if You Did Act. This calls upon (perhaps implicit) knowledge about counterfactual conditionals, and to be effective, the agent's mental state in this case must be sensitive to the state the world is in rather than vice-versa - hence it has a belief-like direction of fit.

Noë, as we will see, is sometimes explicit that Sensorimotor Understanding is not desire-like but belief-like. But to begin, let's review the textual evidence that Noë's account is committed to the idea that Sensorimotor Understanding is desire-like. Although he insistently denies charges of 'behaviourism', he endorses the desire-like direction of fit option at the very beginning of his book, where he says: "Perceptual experience acquires content thanks to our possession of bodily skills. *What we perceive* is determined by *what we do* (or what we know how to do); it is determined by what we are *ready* to do" (2004, p. 1). Throughout the book, he makes the recurrent claim that perception is constituted by the exercise of 'sensorimotor skills' (e.g., Noë, 2004, p. 90), a phrase which without some gymnastics is hard to think of as referring to anything other than capacities of one kind or

another for skillful bodily behaviour. Noë also equates sensorimotor skills with 'practical knowledge', and this, he makes a point of arguing, conforms to Ryle's (1949) view of knowing-how, which Noë suggests means it is an *ability* and does not reduce to a variety of knowing-that.

Replying to Stanley and Williamson's (2001) argument that knowing-how cannot be an ability because you can know how to do something without having the ability to do it, Noë argues that it is in keeping with ordinary language, which he takes Stanley and Williamson to be motivated by, that a ski-instructor who cannot perform the moves they teach, or a pianist who loses the use of their hands, does *not* in fact know how to perform the relevant tasks. Again, if Sensorimotor Understanding is a practical ability and not a form of knowing-that - and if it resembles the pianist and ski-instructor examples - then we must think of it as an ability to make skillful bodily movements. It is therefore desire-like; consider the incoherence of saying that Sensorimotor Understanding involves 'knowing-how what would happen if you did act'.

Noë goes on to contrast Sensorimotor Understanding - which does not necessarily draw on propositional knowledge - with linguistic ability, which does. He claims that a better analogue for Sensorimotor Understanding than linguistic ability is the ability to communicate using bodily gestures (2004, p. 90), although he does not make it explicit what one should take from this analogy. Taken in a minimalistic way, the lesson is merely that neither sensorimotor knowledge nor the ability to gesture call upon a capacity for propositionally-structured thought like linguistic ability does. Noë draws a more direct connection between Sensorimotor Understanding and gesture elsewhere in the book, however, when he suggests that sensorimotor knowledge is actually constituted, at least sometimes, by an ability to gesture. Gesture makes an appearance, in particular, in Noë's account of the skills that underlie the perceptual experience of P-properties. He claims that the ability to experience P-size - the apparent size on an object in your visual field - is equivalent to the ability, for example, to hold out an arm and represent the perspectival size using thumb and forefinger. Endorsing an explicit appeal to capacities for bodily movement made by Pettit (2003), he suggests that experiencing the motion of a ball as it flies through the air involves being able to make bodily movements that manifest a discrimination of its speed, for example by catching

the ball.

Following Evans (1982), Noë suggests that experiencing an object's location in egocentric space, for example its being 'off to the left', is identical to practically understanding the movements you would have to make to face the object head-on. Evans (1985) pointed out that it will not do to specify the relevant movements solely in terms of muscle movements, because the appropriate behaviours could include, for example walking, crawling, swimming and indefinitely many others. Evans concluded from this that we cannot derive the notion of space by specifying appropriate behavioural capacities, but must presuppose space in order to specify what the relevant behavioural capacities are, thus endorsing a notion of 'behavioural space'. In contrast, Noë endorses a notion of "egocentric, sensorimotor space" (2004, p. 89). The concepts of behavioural space and sensorimotor space appear to be the same, or closely related: both characterise the environment spatially and in egocentric reference to the subject's body. Noë does not explain what difference there is, if any between the concepts, and what he probably intends is merely to underline his difference with Evans on the subject of how perceivers master that space.

Here Noë distances himself from the idea that sensorimotor mastery is Being Able To Act. He says:

Evans (1982) sometimes seems to offer such a behaviorist account, suggesting that the experience of something as off to the left consists, as it were by definition, precisely in the possession of certain dispositions to move with respect to the thing. My claim is not behaviorist in the way Evans's appear to be. When we see a flicker on the right we know - in a practical, implicit way - that movements of the eyes to the right bring (or would bring) the flicker better into view (Noë, 2004, p. 89).

This implies that Sensorimotor Understanding is Knowing What Would Happen If You Did Act, which is to say has a belief-like direction of fit. That Noë intends Sensorimotor Understanding to have a belief-like direction of fit is underlined in an exchange between Noë and Campbell concerning the experience of P-properties. Campbell (2008) observes that Noë's account lacks an explanation of how we locate objects in sensorimotor space to begin with. Campbell's point is that Sensorimotor Understanding, taken as an implicit, belief-like mastery of how P-properties change in line with movement, is sufficient to explain the

experience of objective properties, but that we must appeal to a distinct faculty to account for the experience of P-properties. He complains that in Noë's account: "there is no attempt at all to explain what it is to perceive an aspect of an object" (p. 669). Campbell's friendly suggestion is that the sensorimotor account should explain P-property experience by appeal to neurally-realised capacities for actions such as reaching out to grasp the object in front of you. This suggestion does not appear to be very far removed from the spirit of Noë's proposal. Noë (2008) rejects this, however, on the ground that that perceptual experience is not *for* action. He responds that experiencing P-properties implies knowing what movements one *could* make to gesture to the object or to bring it into view, but emphasises that this knowledge does not require an ability or disposition to make the relevant movements.

As an aside, Noë's reply to Campbell is unconvincing. He implies that as long as we do not stipulate that P-property experience happens *prior* to objective-property experience, we need not be committed to the view that they are the result of separate faculties. But the import of Campbell's point, concerning the 'to the left' case for instance, is that in order to know which movements would cause you to face the object head-on, you need some way to recognise the state of affairs that obtains when you are facing the object head-on. You may know that *no* movements are required to face the object head-on, and that you are therefore facing the object head-on. But in this event, some account must be given to explain how you know that no movements are required. Adopting Campbell's proposal, we could propose that you know this because you have the ability to carry out a practical task like picking the object up, but Noë rejects this solution. I will examine shortly how this limitation in Noë's account might be redressed while maintaining that Sensorimotor Understanding has a belief-like and not merely a desire-like direction of fit.

Elsewhere, Noë (2004) repeatedly disavows 'behaviourism' on various grounds. He states that it is a mistake to suppose that "effects are logical constructions of their causes" (p. 118), i.e. that the behaviours that are the effects of Sensorimotor Understanding themselves constitute Sensorimotor Understanding. He claims that Sensorimotor Understanding is not a set of behavioural dispositions, but the *ground* of your dispositions (p. 88). Importantly, he notes that you do not always need to move, or even have the ability to move, in order to

perceive (p. 90).¹⁴ These are sound objections to certain positions that might be deemed behaviourist, and I will later attempt to demonstrate that a position which does claim that Sensorimotor Understanding depends on behavioural capacities can be finessed in a way that defuses these worries (sections 4.4 and 6.2).

The uncertainty around direction of fit leaves us with two opposed possibilities. On a belief-like account, Sensorimotor Understanding is a non-propositional belief-like understanding of the sensory consequences of possible movements. This possibility is difficult to square with Noë's claims that the understanding is practical knowledge or a set of sensorimotor skills, and with his claim that perceptual consciousness depends on what we are able to do. But it does justice to Noë's denial that Sensorimotor Understanding is for action, or depends on movement or the ability to move. Accounting for the physical processes that enable Sensorimotor Understanding, on this view, appears to require explaining how neural states and processes yield belief-like states, and this may require endorsing one kind of internal representation account or another, since it is not apparent what other kinds of subpersonal state could do the necessary explanatory work.

Sensorimotor Understanding, on a desire-like account, could involve knowing what movements to make to achieve various practical goals, and this knowledge might itself depend on the deployment of action-oriented representations. Alternatively, Sensorimotor Understanding may, on the desire-like account, consist of *dispositions* or *capacities* for various behaviours which are not best accounted for by reference to representation or 'knowledge' at all. Noë, at times, resists both these possibilities. In section 4.4, below, I will argue that one can avoid the most important worries associated with behaviourism while nonetheless taking behavioural capacities to be necessary and sufficient for the existence of Sensorimotor Understanding. To this effect, I will attempt to establish that Sensorimotor Understanding exists by virtue of the subject's possessing an appropriate set of behavioural capacities, but that it 'grounds' the capacities rather than being identical to them. This thesis, as will become evident, is compatible with the further claim that representations are required to explain how

¹⁴ Noë's other objection is that behaviourism denies the existence of phenomenal consciousness (2004, p. 32). However, it is not obvious that explaining experience by appeal to behavioural dispositions necessarily entails denying consciousness.

the relevant behavioural capacities are realised, but also with the claim that nonrepresentational explanations are more appropriate. Before I get to this, I will examine a more straightforwardly representationalist option, involving the claim that Sensorimotor Understanding is Knowing What Would Happen If You Did Act, and depends on the deployment of a particular kind of internal representation.

4.3 - Sensorimotor Understanding as Subpersonal Representation

Suppose, in accordance with some of Noë's (2004) claims, that possessing and exercising Sensorimotor Understanding neither requires bodily movement nor an ability or disposition to move. To adopt this position is, by definition, to reject a number of claims that could be deemed behaviourist. In addition, suppose that action is not what Sensorimotor Understanding is primarily for. Instead, let's take it that Sensorimotor Understanding is a kind of knowledge about the P-properties that presently obtain, and the P-properties that would obtain given various possible movements. As I put it earlier, Sensorimotor Understanding, on this view, has primarily a belief-like direction of fit. Sensorimotor Understanding continues to have one non-trivial conceptual link with bodily behaviour, namely that it incorporates knowledge about the consequences of possible bodily movements.

Noë (2010) continues to deny that conscious vision is primarily for action, but argues that if action-guidance is one purpose of conscious vision, then conscious vision requires Sensorimotor Understanding. He contrasts the sensorimotor approach, as usual, with the claim that perception consists of the activation of an inner world model. On the inner model conception, perceiving involves knowing what objects are out there, but not what bodily relation you stand in to them. Lest one think that Noë is targeting only a pictorial theory which no one endorses, notice that any representationalist theory is in this position by default if the representational content it endorses does not somehow specify a bodily relation. Noë claims that if your perceptual awareness did not represent your bodily relation to the world, then the world would not seem phenomenally 'present' to you. Even if this is not true, it is not apparent how your perceptual awareness in this event could help guide action. Noë observes that it must be possible to integrate what you see with what you do, and argues that it is a

virtue of the sensorimotor theory that it shows how this integration is possible.

An upshot of this point, we might think, is that Sensorimotor Understanding necessarily involves an ability to make skillful bodily movements, so long as something else does not block the ability - for example paralysis, or optic ataxia, an inability to co-ordinate muscle movements. This is one way of weakly motivating the conceptual link between perception and action implied by the title of Noë's (2004) book. It also justifies the description of Sensorimotor Understanding as practical, but only if we take 'practical' to mean knowing-that *and* knowing-how, as Stanley and Williamson do and Noë denies. Sensorimotor Understanding, on the present view, is knowing what would happen if you moved, a knowing-that, accompanied by a knowing-how to act, which depends on the knowing-that. However, the present link between perception and action is less tight than the one expressed by Noë's claim that "*What we perceive* is determined by *what we do* (or what we know how to do)" (2004, p. 1), since the present view is in essence only signed up to the more pedestrian claim that what we do or know how to do is determined by what we perceive.

Contradicting Noë's (2004) discussion of practical knowledge, Sensorimotor Understanding on the present view is not merely a practical ability, but a belief-like understanding of SMCs that enables behaviour. This means that the sensorimotor approach owes an explanation of what Sensorimotor Understanding is and, in particular, how it is physically constituted or realised. Since it is belief-like, it appears to require truth-evaluable content. Despite asserting that we cannot usually express our understanding of SMCs using propositions or bring it to bear in "explicit deliberative judgement" (p. 187), Noë asserts that Sensorimotor Understanding is conceptual, on what he takes to be a legitimately unrestricted account of what conceptual thought requires. He changes this position in later work, asserting instead that Sensorimotor Understanding involves a kind of nonconceptual content, but maintaining that Sensorimotor Understanding occurs on a continuum with conceptual skills rather than belonging to a radically different kind (Noë, 2012). This does better justice to the suggestion that Sensorimotor Understanding involves content which the perceiver cannot fully express using propositions, even, presumably, by pointing to an object and saying '*that* sensorimotor law'. This raises the possibility that Sensorimotor Understanding supervenes on the deployment of neural representations bearing nonconceptual content.

Seth (2014) has recently offered a specific representationalist proposal intended to do justice to some aspects of the sensorimotor theory. His account rejects the sensorimotor theory's explicitly enactivist apparatus, namely anti-representationalism, vehicle externalism and the endorsement of dynamical systems explanations. However, he endorses the account's approach to naturalising phenomenal qualities, and, in particular, Noë's (2004) insights into visual phenomenology, including presence-in-absence. Seth combines these aspects of the sensorimotor theory with predictive coding (see, Clark, 2013; Hohwy, 2013). Like the orthodox approach, predictive coding accounts of perception hold that the brain constructs internal representations designed to represent features present in the outside environment as accurately as possible, based on the limited information available at one time via the sense organs. Orthodox theories make sensory inputs the starting point, and propose that the brain uses them to build up an inner model. By contrast, predictive coding accounts hold that the brain constructs models of the environment which it 'predicts', using Bayesian inference, will prove to be in keeping with the worldly state of affairs that really obtains.

The difference between this and the orthodox approach is that sensory inputs here serve as error correction signals; instead of acting as precursors to the construction of an inner model, the sense inputs function to tell the brain when the inner models it has already deployed are wrong. Further, the predictive approach holds that the brain simultaneously holds multiple inner models organised hierarchically into tiers at increasing levels of abstraction. Only the lowest model in the hierarchy uses sense inputs themselves as error correction signals. This model functions as an error correcting signal for the model above, and so forth. An example of a substantive explanatory difference between orthodox and predictive coding accounts occurs in the phenomenon of *repetition suppression*, a tendency of the brain to respond less to a stimulus as it becomes more familiar. Seth cites work showing that repetition suppression fails to occur when the subject does not expect the stimulus to be repeated (Summerfield, Trittschuh, Monti, Mesulam and Egner, 2008). He observes that this is intelligible within a predictive coding approach, since the approach makes expectations a central feature, whereas there is no natural way to make this intelligible within an orthodox framework, where sense inputs come before the models which make use of them.

Seth endorses some of the Noë's (2004; 2012) insights about visual phenomenology.

One example of this, which I have discussed already, concerns shape-related P-properties and objective properties. You are confronted with a coin, tilted at an angle, which is therefore elliptical in aspect. Thanks to your Sensorimotor Understanding, you appreciate both that it has an elliptical aspect and that objectively it is circular. It is a key virtue of the sensorimotor approach that it explains how you are sensitive to objective properties while at the same time, necessarily, being sensitive to them from some perspective rather than from nowhere. Another aspect of this phenomenon is ‘presence-in-absence’, the tendency of environmental features which are strictly speaking out of view to show up in visual consciousness anyway. Noë’s (2004) prototypic example concerns looking at a tomato. In one sense, you can never see the whole tomato at once, but only one side of it, while the other side remains out of view. According to Noë’s phenomenology, which Seth endorses, the back of the tomato nonetheless has presence-in-absence, which is to say your visual experience is of a whole tomato rather than something like a 2D cut-out.

Seth argues that a predictive coding account can do justice to this, proposing that the predictive models employed by the brain in perception incorporate not only predictions about the features of the present environment causally responsible for present sense inputs, but also counterfactual predictions about the changes in sense input that would result from a range of possible movements. Perceptual presence is explained by these counterfactual predictions, and the more predictions there are about movement-related counterfactuals relating to some feature, the greater the degree of felt phenomenal presence there is associated with that feature.

For my present purpose, the crucial feature of Seth’s account is that it provides a better developed explanation of the mechanism that underlies Sensorimotor Understanding than found in the original statements of the sensorimotor approach, and does justice to the ‘belief-like’ construal of Sensorimotor Understanding I have been addressing in this section. It is notable that adopting this view means adopting a very conservative rendering of the sensorimotor account. It means endorsing representationalism, rejecting dynamical entanglement and ECM, and by focusing on internal models, downplaying the role played by embodied interaction even as causal scaffolding.

Sensorimotor mastery, and in particular Sensorimotor Understanding, one of the

most important notions in the sensorimotor theory, is ambiguously defined by the theory's original statements. At times, it appears to be implicit knowledge of the SMCs that presently obtain, and nothing more. In this case, the obvious way to explain Sensorimotor Understanding is in keeping with orthodox cognitive science via a representationalist story. Seth's account has the virtue of doing justice to Noë's (2004) insights about visual phenomenology and perceptual presence, and of allying this aspect of the sensorimotor theory with insights from the currently popular predictive coding approach. If nothing else, Seth's account also has the virtue of beginning to repair the sensorimotor theory's serious explanatory deficit by accounting for the precise mechanism that realises Sensorimotor Understanding. However, this comes at a cost. For anyone who is motivated, on whatever independent grounds, to endorse the enactivist theses of anti-representationalism, vehicle externalism and dynamical systems, Seth's account will not help, since it serves only to co-opt insights from the sensorimotor theory into a much more conservative representationalist framework. In particular, accounting naturalistically for representational content brings with it a heavy and as yet unresolved explanatory burden (Hutto and Myin, 2012), as I discuss further in chapter 6.

Moreover, the approach has costs for the sensorimotor theory itself. If we endorse internalism and representationalism, then the sensorimotor theory no longer gets to make distinctive use of the appealing proposition that the brain offloads the burden of modelling onto the environment itself. Although any variety of representationalism can observe that the amount represented at any one time is minimal, and that as active perceivers we move around to acquire new information, the implication of Noë's (2004) account is that Sensorimotor Understanding is computationally more efficient than inner modelling, because we merely need to be poised to interact with the environment, rather than modelling the environment and *then* becoming poised to interact with it on the basis of this model.¹⁵ Seth's proposal also has the drawback, from the perspective of sensorimotor theory, of making the sensorimotor account hostage to phenomenological fortune, hanging as it does in large part on Noë's

¹⁵ Of course, whether the non-representationalist version of the sensorimotor account really implies a more efficient process than a representationalist alternative depends on how the representationalist and non-representationalist alternatives are precisely cashed out.

description of visual phenomenology. Clark (2008), for instance, suggests that the feeling of presence-in-absence could be well accounted for by the general feeling that there is more to be seen, rather than any more explicit representation of or knowledge about the sensory consequences of possible movements. For these reasons, I endorse a different view of Sensorimotor Understanding, which I lay out in the next section.

4.4 - Sensorimotor Understanding in Bodily Capacities

Notwithstanding suggestions by various commentators that the sensorimotor account requires representation, the theory is, in at least one sense, strongly anti-representationalist. The idea is that there is no need to represent the world, because it is already present to you. This is reflected by the outside memory thesis (O'Regan, 1992; O&N), and in Noë (2004; 2012) by the idea of presence-as-access, which implies that perception does not depend on representing even a small part of the outside world. One way of thinking about the theory's anti-representationalism is to think of it as a constitutive claim about perception at the personal level. On this proposal, we are not, as persons, in touch with representations of the world, but stand in an unmediated relation to the world itself. This is not necessarily incompatible with the claim that perception depends on internal representations subpersonally, a point I return to later.

The sensorimotor account's way of cashing out this (personal level) relation with the outside world is to say that it involves possessing and exercising a set of bodily capacities; specifically, capacities to skillfully interact with the outside environment in ways that draw on an implicit understanding of the SMCs that obtain. Consider that on a representationalist account, perception involves representing the world, and interacting with the environment in a goal-directed way is a contingent consequence of the fact you perceive it. Since the role of the 'representation' is played in the sensorimotor account by the world itself, it follows that perceiving itself in the sensorimotor account is equivalent to possessing and exercising the capacity to skillfully interact with the world in a goal-directed way. This provides a solid motivation for the claim that perception depends on action and Being Able to Act, which, as I said earlier, implies a desire-like direction of fit.

The question remains, however, of how we account for Sensorimotor Understanding in its guise as a belief-like faculty, a Knowing What Would Happen If You Did Act. The best approach, I suggest, is to avoid reifying Sensorimotor Understanding. We should not think of it as the sort of thing that depends, by definition, on a particular sort of concrete entity, for example a conscious or purely subpersonal representation, which causally mediates behaviour. Instead, we should think of it as a set of necessary conditions that our capacities for practically-oriented, goal-directed behaviour must meet to implicate perception and perceptual consciousness.

To understand the explanatory relation between Sensorimotor Understanding and behaviour, it will help to distinguish between some different varieties of explanatory relation theoretical entities can stand in to one another. Bermudez (2000) distinguishes between ‘vertical’ and ‘horizontal’ explanations in cognitive science (not to be confused with Hurley’s 1998 usage of the vertical/horizontal metaphor). Horizontal explanations exist at one level, the personal or subpersonal, and describe temporally distributed events. So if we are offering a horizontal explanation and we say that A causes B, we mean that A happens and this causes B to happen a short time later, where A and B occur at the same explanatory level. On Seth’s (2014) proposal, we might imagine that neural representations of the sensory results of possible movements send information to motor output areas of the brain, and so generate behaviour. This would be a horizontal explanation. Vertical explanations describe events that happen simultaneously at different explanatory levels. In this case, A might be a subpersonal computation, and B might be a psychological state possessed by the whole person. People sometimes use the word ‘cause’ in vertical explanations too, although it is not causation of the sort that implies temporal extension. We might also say, here, that A *enables* B. If we said that the deployment of the subpersonal representations described by Seth enable perceptual consciousness, this would be a vertical explanation.

The important thing, for the moment, is that the explanatory relation between Sensorimotor Understanding and a capacity for skillful behaviour is neither a vertical nor a horizontal one. Sensorimotor Understanding, I propose, is a logical condition that the behaviour enabled by our behavioural capacities must meet in order to constitute perception and perceptual consciousness. The identity between perception and skillful behaviour arises

from the fact that to perceive is not to model the world and then use this model, but merely use the world in the way traditional theories hold that we use an inner model. The fact that the behaviour needs to manifest a belief-like Sensorimotor Understanding is an upshot of the straightforward claim that it would be impossible to skillfully interact with the world in more than a very basic way if you did not have the ability to anticipate, in at least some cases, the sensory results of your movements. On this basis, behavioural capacities manifesting Sensorimotor Understanding can be described as a condition of possibility for perception. This is one way of understanding the claim that Sensorimotor Understanding *grounds* perception-implicating behaviour.

The view I am advocating is neutral with regard to whether or not Sensorimotor Understanding is enabled by the deployment of subpersonal representations. However, it does not require internal representation as a matter of definition or logical necessity. Sensorimotor Understanding should be thought of as a characteristic of behaviour; an adverb or adjective used to describe behaviour or a behavioural capacity, and not a noun. Although there is no adverb ‘understandingly’, we might say that a behaviour ‘shows understanding’, and mean that the understanding is constituted by the behaviour rather than that the behaviour merely demonstrates the existence of a pre-existing entity called understanding. Since Sensorimotor Understanding is a logical condition for certain kinds of skillful behaviour, appropriate behaviour is almost sufficient by itself to guarantee the truth of an ascription of Sensorimotor Understanding. We should add to this an extra requirement that the brain makes an intelligible contribution to the behaviour, since it is doubtful that a person in a metaphysically possible world who exhibits appropriate behaviour but lacks internal states to enable the behaviour could be conscious.

It may help to compare the current view of Sensorimotor Understanding with Dennett’s (1987) intentional stance, which I discussed in chapter 2. According to the intentional stance, mental states such as beliefs and desires exist just as long as ascribing them to a person or other agent helps an outside observer explain the person’s behaviour. Like Sensorimotor Understanding, this implies the existence of a set of necessary and sufficient conditions that an agent’s behavioural repertoire must meet for the agent to possess a particular mental state. Sensorimotor Understanding is correctly ascribed to any perceiver

whose behavioural repertoire meets the appropriate conditions. In contrast with a certain understanding of the intentional stance, however, Sensorimotor Understanding is, in a robust sense, real, and not observer-relative. This is because certain kinds of skillful behaviour require, as a matter of necessity, a sensitivity to certain presently obtaining SMCs.

If we can account physicalistically for the relevant behavioural capacities without appeal to neural representations, this will be enough to account for perception, and the belief-like Sensorimotor Understanding will come for free. The suggestion that it is possible to explain the relevant behavioural capacities without appeal to representations is controversial, of course, and nothing I have said so far shows this can be done. However, it is important that Sensorimotor Understanding, on the view I propose, does not require subpersonal representation by conceptual necessity. Moreover, the present view is compatible with the possibility that perception does involve neural representations, but not neural representations of SMCs. This holds open the possibility of a subpersonal representationalist explanation of Sensorimotor Understanding quite different from Seth's approach.

We need, of course, to know something about the behaviour required to justify ascriptions of Sensorimotor Understanding. In the case of P-properties like relative location, P-size and P-shape, Sensorimotor Understanding is manifest in the capacity to make appropriate movements in sensorimotor space. For instance, an understanding that something is to the left might be manifest in the capacities to walk, run, crawl (etc.) in a leftward direction for the purpose of bringing the object into view. When we experience something as being to the left, this usually implicates, as Evans (1985) says, a capacity to walk *or* run *or* crawl, or perform any action from an indefinitely wide range of possible actions in order to reach the object. If we know that someone intends to face the object head-on, and we see them walking toward it, we could, in a manner reminiscent of the intentional stance, interpret them as possessing an implicit understanding of the fact that this movement will bring the object into view, hence as possessing Sensorimotor Understanding. However, ascribing a perceptual state to a person or animal on the basis of this single behaviour would be ascribing too much. The behaviour could be a reflexive response to a stimulus, where perception ought to enable an ability to respond flexibly using a wide repertoire of behaviours. So the subject ought to have the ability to run, crawl, etc., as well as walk in order to reach the stimulus; the capacity

has, in this sense, to admit some degree of generalisability. We can do justice to Noë's contention that the experience of objective properties is grounded by the same faculty of Sensorimotor Understanding as the experience of P-property if we cast the experience of objective properties as involving the exercise of a behavioural capacity, too. In common with Sensorimotor Attunement, appealing to this capacity differs from the behaviourist appeal to input-output relations, because it does not specify which outputs should follow when particular inputs obtain, but which outputs should follow when particular SMCs obtain.

There are various similarities and differences between Sensorimotor Understanding and the notion of Sensorimotor Attunement, which I reconstructed from O&N in section 3.3.1. In Noë (2004), Sensorimotor Understanding is largely if not entirely framed as a personal-level capacity, meaning that its operations are accessible to the conscious subject. This reflects the fact that we are conscious of possessing and exercising the relevant understanding, or at least we always have the possibility of becoming conscious of it, even if there are cases in which we are not conscious of our behaviour because we are sufficiently absorbed. Evans (1985) claims that to experience an object as occupying a particular spatial location, the conscious subject must understand the movements required to get there. This might be equivalent to nothing more than being aware of what your bodily dispositions are. If you are aware of the movements you must make to reach an object that is to your left, for instance, then you are aware of the SMCs associated with the object's being on the left, i.e. the movements that will result in your facing the object head-on. In this sense, the SMCs involved in 'being to the left' are transparent to you (i.e. you know what they are). The experience of P-colour, the colour a surface presents when viewed from a particular location and under a particular lighting condition, is a different sort of case. From the perspective of a conscious perceiver, there is no transparent connection between this aspect of colour phenomenology and possibilities for bodily movement. Noë (2004) thus suggests that P-colour experience is accounted for by a subpersonal (for which we can read subdoxastic) variety of sensorimotor mastery (p. 140). Since I am attempting to offer a maximally coherent reconstruction of the sensorimotor theory, I will take it that this is accounted for by Sensorimotor Attunement.

This brings to the fore an important difference between Sensorimotor Attunement and Sensorimotor Understanding. When I attempted to do justice to O&N's notion of

Sensorimotor Attunement, I noted that there was more than one way the theory could propose we identify the specific SMCs to which a perceiver is presently attuned when she exercises the relevant behavioural capacity. I suggested that one option is to endorse a Dennett-style interpretative stance and identify the SMCs by appeal to a combination of teleology or proximal goal and present behaviour. I went on to endorse a more austere alternative in which we suppose that a goal-directed agent masters a given set of SMCs if she responds differentially to them in her behaviour, regardless of whether or not one could identify what the SMCs are by running the interpretative stance.

In this event, the agent does not have to know or represent what the SMCs are, she merely has to discriminate them. This is compatible with the suggestion that conscious perceivers do not have cognitive access to the SMCs involved in P-colour. This contrasts, however, with the conscious Sensorimotor Understanding that applies to P-shape, location and movement. In this case, where we do have the option of running a behavioural test to see if a subject possesses Sensorimotor Understanding, the subject in some sense consciously knows what the SMCs are, because the SMCs themselves, and not just discrimination of the SMCs, are identifiable from the subject's behavioural repertoire - e.g., she is going to the left, she wants to face the object head on, so she has mastered the fact that going to the left will result in facing the object head on.

To conclude this section, I will attend to some interesting remarks by Glock (2008) on the subject of animal thought. In the piece, Glock attempts to tread a middle path between linguisticism, which claims that non-human animals lack thought altogether because they lack language, and the mentalist claim that animal thought differs from human thought only in degree and not kind. Although this is a quite different topic to the one that concerns us here, it chimes with Noë's (2012) stated ambition of accounting for continuities between thought in humans and other animals, and between non-linguistic thought and sensorimotor skills in mature humans. Glock criticises Searle (1994), who claims that the best way to identify whether or not an animal is conscious is to identify neurophysiological features correlated with consciousness and check whether the animal possesses them or not. Glock points out that this only works if you have a way to test outwardly observable behaviour for consciousness in the first place. Even once you have identified the relevant features in one

species, there is no guarantee that another species which also possesses those features is also conscious purely because it also possesses them. Glock's solution to the question of animal minds is to focus on capacities that can be manifested in behaviour instead. The capacities in question are not mere dispositions to discriminate one object from another, but more flexible abilities, for example, to respond to distinct objects differently in one situation but not another, and to correct mistakes. The idea is not that thoughts reduce to certain behavioural dispositions, but that they ought only to be ascribed on the basis of them.

In the case of human perceptual consciousness, the problem is different. There is no need to concern ourselves with the question of whether or not other humans enjoy perceptual consciousness. But a question does arise with regard to what perception consists of. It is possible that the only thing which would licence us to claim that a particular neurophysiological state gives rise to perceptual consciousness is its being correlated with a behaviour too. As I have suggested here, that is behaviour meeting a certain set of conditions described as Sensorimotor Understanding. This is an epistemic motivation for ascribing perceptual skills on the basis of behavioural capacities.

However, there is a deeper motivation, which is the possibility that perception just *is* a capacity for skillful bodily interaction with the environment. There is no default reason for thinking this claim fares any worse than the claim that perception is constituted by the activation of an inner world model. If the capacity claim is true, we get the benefit of a direct realist account of phenomenal qualities, which I mentioned in section 3.6 in the discussion of ECM, as well as the epistemic benefit of a non-mysterious way to explain how physical states and processes give rise to perceptual consciousness. In the next chapter, I claim that sensorimotor theory can enjoy partial success as an account of phenomenal qualities even if the present view fails. In chapter 6, I defend the claim that perception is constituted by the possession and exercise of a set of behavioural capacities.

5 - Phenomenal Character: A Sensorimotor-Representation Hybrid

The next two chapters address a response made by Clark (2008) to the sensorimotor account in which he claims that perception plausibly consists of the construction and deployment of representations standing in for types, categories and relative locations, and not the exercise of sensorimotor skill. In chapter 6, I pursue some more radical claims arising from the view of Sensorimotor Understanding as a set of conditions that behaviour must meet, and the attendant possibility that perception does not involve representation at all. In this chapter, however, I explore a response by the sensorimotor theory to a worst case scenario in which it is established that perception does depend on representation, and not on behavioural capacities. I argue that even in this event, the sensorimotor theory has the ability to provide a tenable, compelling account of phenomenal qualities. This chapter is structured as follows. In the first section, I outline a criticism levelled by Clark (2008) against the sensorimotor theory. In the second section, I propose a line of response that could be adopted by a conservative rendering of the sensorimotor approach. In the third section, I examine an apparent problem for this response, and explain how it can be overcome. In the fourth section, I offer further reasons to endorse the sensorimotor account of phenomenal qualities. In the fifth section, I examine the relation between sensorimotor theory and Rupert's (2011) model of the 'massively representational mind'.

5.1 - The Non-Sensorimotor, Skill-Based Approach

In a critical response to the sensorimotor theory, Clark (2008) endorses, in broad outline, an approach to perceptual consciousness which he suggests is both more plausible than the sensorimotor account, and incompatible with its key claims. Clark's preferred approach in part comprises a philosophical approach to phenomenal qualities or 'qualia', which he argues are properly accounted for by the perceiver's exercise of 'epistemic' skills and not the sensorimotor skills that Noë and O'Regan appeal to. The approach also comprises an empirically-oriented account of the cognitive processing that underlies perceptual consciousness. Drawing in particular on Milner and Goodale's work on dorsal and ventral

stream perception (e.g., Milner and Goodale, 1995), Clark suggests that the empirical evidence demonstrates that conscious perception is not primarily the result of a sensorimotor process. I will begin by outlining Clark's preferred approach to phenomenal qualities, and subsequently move on to his claims about cognitive processing.

Skill-based accounts of phenomenal quality propose a way to account naturalistically for the phenomenal qualities featured in perceptual consciousness, for example the look of red. On this approach, a given phenomenal quality is accounted for by nothing more or less than the subject's ability to skillfully discriminate, in some relevant way, a feature of the environment associated with the quality from other features they have the ability to consciously perceive. One such account is offered by Pettit (2003). As Pettit notes, traditional accounts imply that a subject's capacity to discriminate one property from another is merely the contingent outcome of a difference in phenomenal look. Conversely, skill-based accounts claim that there is nothing more to experiencing a given phenomenal look than exercising the associated discriminatory capacity. In Pettit's examples, you exercise these capacities in making judgements, for instance that an object is red, or a brighter shade of red than some other object, as well as in bodily behaviour, for example reaching out appropriately to catch a fast-moving ball. An object looking a particular shade of red to you is, on this view, no more than it looking to you like you can discriminate it in an appropriate way from other shades of red, and other colours, and so forth. Presuming that there is no fundamental obstacle to describing the exercise of these capacities for discrimination in physicalistic vocabulary, there is no deep-seated obstacle on this approach to understanding phenomenal qualities physicalistically.

Clark takes it that the sensorimotor theory is itself a skill-based approach. On most ways of understanding the account this is true, although we must take care to be clear on why this is so. The sensorimotor approach proposes that the look of red, like other phenomenal qualities, is properly accounted for by a sensorimotor law, which is a systematic way that sense input signals originating from a relevant object change in line with movements and varying ambient lighting conditions (e.g., Noë, 2004). One way to understand the concept of sensorimotor law, therefore, is to think of it as something that can be described solely by reference to counterfactual conditionals of the form 'if the subject moved like this, sense input

would change like this' (suggested by Noë, 2004). So understood, sensorimotor laws comprise what Burhmann, Di Paolo and Barandiaran (2013) describe as the subject's *sensorimotor environment*. These laws have a motor output to sense input direction of fit, since they characterise how sense input will change in line with given movements, but do not describe how the subject will move given particular sense inputs. So understood, sensorimotor laws may, in principle, be sufficient to explain degrees of similarity and difference between specific phenomenal qualities where it is assumed that the subject is in some appropriate way sensitive to those sensorimotor laws. But they could not explain how a perceiver can be in conscious states characterised by those qualities, considering that the relevant counterfactual conditionals might be true even if the perceiver was dead, or their brain were made of jelly (etc.).

To explain how perceivers do come to be in the relevant conscious states, the sensorimotor approach posits that in some way or another we master sensorimotor laws. Depending on how we construe the notion, mastery of the sensorimotor environment could include internal representations of counterfactual conditionals, perhaps even propositional knowledge, without necessarily calling on any bodily skills whatsoever. As it happens, however, most versions of the sensorimotor account individuate phenomenal qualities by reference only to patterns of counterfactual dependence that the perceiver manifests a sensitivity to in their bodily movements.¹⁶ Burhmann, Di Paolo and Barandiaran (2013) characterise the perceiver's repertoire of bodily responses in a given sensorimotor environment through what they term a *sensorimotor habitat*, which this time comprises the movements the perceiver makes in line with given sense inputs and the changes in sense input that follow. The need to refer to a perceiver's sensorimotor habitat and not just their sensorimotor environment is evident in O&N and O'Regan (2011), regardless of their further appeal to higher-order

¹⁶ Noë (2010) actually denies that the sensorimotor theory is a skill-based approach, on the ground that perception depends on understanding SMCs rather than merely being able to skillfully discriminate them. However, the frequent invocation of sensorimotor 'skill' invites a reading in which the sensorimotor theory is a skill-based approach. This reflects the tension between competing directions of fit I discussed in section 4.2. Since I concluded that Sensorimotor Understanding ought to be thought of as depending on knowing how to act as well as knowing what would happen given possible movements, I think it appropriate to consider the sensorimotor theory a skill-based approach.

tracking of those bodily interactions. Noë (2004) does not say anything explicit to suggest that a perceiver can only be conscious of sensorimotor laws featured in their sensorimotor habitat. However, the claim is not explicitly retracted in Noë's (2004) solo account. At times, Noë attempts to avoid endorsing an unduly strong link between perception and action, and in particular denies that you necessarily have to be moving your body right now in order to see. This does not compromise the suggestion that you can only have perceptual experiences characterised by SMCs featured in your sensorimotor habitat, so long as we take the sensorimotor habitat only to characterise ways you are *disposed* to skillfully interact.

Clark agrees that phenomenal qualities should be explained by reference to the exercise of skills. However, instead of appealing to bodily responsiveness to patterns of sensorimotor dependence, his favoured account appeals to personal-level responsiveness in reasoning and action planning to categories, types and locations. He takes it that 'epistemic' skills of this sort either do not call upon capacities for bodily movement, or at least do not do solely call upon such capacities, and are distinct from the bodily skills appealed to by the sensorimotor theory. To keep things unambiguous, I will describe epistemic skills, as Clark understands them, as 'non-sensorimotor' skills, leaving the designator 'epistemic' neutral on whether or not the skills in question are identical to capacities for bodily movement. Importantly, Clark argues that sensorimotor and non-sensorimotor skill-based accounts of phenomenal qualities share a common explanatory virtue when it comes to accounting for phenomenal qualities non-mysteriously; indeed, nothing he says is intended to call into doubt the view that the sensorimotor and non-sensorimotor skill-based accounts are, in principle, equally well-suited for resisting the claim that phenomenal qualities cannot be accounted for naturalistically. This makes it plausible that Clark thinks the appeal to skills is enough to solve this problem without even getting into the question of whether they are sensorimotor or not. His objection to sensorimotor theory, by contrast, is that it provides an implausible account of the capacities that are, as a matter of fact, deployed in perceptual consciousness.

Clark's objection partly reflects a complaint that the sensorimotor theory is untenably committed to supposing that the range of phenomenal qualities featured in an agent's perceptual awareness is no less fine-grained than the discriminations evident in the agent's most finely sensitive bodily behaviour. This commitment is a consequence of the

sensorimotor theory's claim that two perceivers could only share identical perceptual phenomenology if they are identically embodied. O&N, for example, say: "For two systems to have the same knowledge of sensorimotor contingencies *all the way down*, they will have to have bodies that are identical *all the way down*" (p. 1015), with a similar claim appearing in Noë (2004).

In fact, it is not apparent that the sensorimotor theory has any essential need to make the claim. Logically-speaking, the sensorimotor version of the skill-based view only needs to say that for every difference in phenomenal quality there is a corresponding difference in the bidirectional patterns of dependence between input and output featured in the perceiver's skill-driven bodily interactions. It need not make the inverse claim that every discrimination manifest in an agent's behaviour must be reflected in the phenomenal character of their perceptual consciousness. The sensorimotor theory faces pressure to say that the input-output patterns relevant to perceptual experience are very fine-grained in order to avoid an objection levelled by Block (2001). Block's charge is that the sensorimotor theory must claim that a rudimentary machine, differently embodied but manifesting comparable input-output relations, undergoes human-style perceptual experience. O&N's reply is that sameness of perceptual experience depends on sameness in fine-grained input-output relations, which themselves depend on being embodied in a particular way. However, they need not have claimed that sameness of embodiment is required *all the way down*. The sensorimotor theory could respond to the charge just as well by claiming that two perceivers must be embodied in a sufficiently *similar* way in order to enjoy identical perceptual experiences.

If Clark's view were merely that we should account for qualia by appealing to coarse-grained bodily skills, it would not be inconsistent with the most indispensable commitments of the sensorimotor theory. More seriously, however, his comments imply that the relevant capacities may not be bodily capacities at all. The capacity to reason, certainly, does not on the face of it appear to be a bodily skill, and nor does the capacity to plan an action if this is taken to mean choosing an action rather than actually performing it. Clark bases his response largely on empirical work, which he alleges shows that perceptual consciousness is not realised by sensorimotor coupling but a process involving the deployment of internal representations that are in large part unaffected by occurrences at the sensorimotor

periphery.

In particular, he cites Milner and Goodale's (1995) dual systems account of vision. According to the dual systems model, the brain's dorsal stream is associated with a visual system responsible for fine-grained sensorimotor coupling but not conscious visual awareness. Clark does not raise any objection to describing this bodily behaviour as the exercise of sensorimotor mastery, and we can plausibly take this activity to be identical to the activity characterised as the lowest tier of sensorimotor mastery in O'Regan's (2011) multi-tier sensorimotor account. The other visual system is realised primarily by activity in the ventral stream, and is alleged by Clark to enable coarser-grained "reasoning, choice and action selection" (Clark, 2008, p. 190), a description which is strikingly akin to the "planning, decisions, or other rational behaviour" (O'Regan, 2011, p. 121) featured in the second of O'Regan's four tiers. Clark argues that ventral stream activity is not sensorimotor mastery but involves the deployment of representations that, in line with the epistemic skill-based approach to qualia, stand-in for attributes such as "categories, types, and relative locations" (p. 192), optimised for the exercise of epistemic skills, i.e. to enable to person to "sift, sort, select, identify, compare, recall, imagine and reason" (p. 192).

In this vein, Clark says: "[T]he best functional and architectural explanation, according to Milner, Goodale, and others, is that conscious perceptual experience reflects the activation of representations that have less to do with the fine details of world-engaging sensorimotor loops and more to do with the need to assign inputs to categories, types, and relative locations" (p. 192). The role of representation here needs to be treated with care. O'Regan (2011) explicitly denies that conscious states can be adequately accounted for by reference to the 'activation' of representations, a locution featured in the Clark quote above, on the ground that accounts of this sort cannot adequately explain why the conscious states in question have one phenomenal character rather than another. Clark appears to endorse O'Regan's view, at least insofar as he endorses a skill-based account of phenomenal qualities and not, for instance, an account that says phenomenal qualities are accounted for directly by the content of the purported representations (cf., Tye, 2000). Being mindful of this, Clark's view ought to agree that it is not *just* the 'activation' of a representation, but the contribution the representation thereby makes to the exercise of the relevant personal-level skillful activities

like sorting and sifting that contributes to an explanation of phenomenal qualities. At the subpersonal explanatory level, it is an open question how a given ventral stream representation contributes to the exercise of personal-level skill. An explanation here may appeal to some combination of its content, its functional role within a given system, and that system's functional role in the wider cognitive economy. Whichever features an endorser of this approach appeals to, they must explain what contribution these features make to the exercise of a personal-level ability to skillfully discriminate aspects of the visual environment.

In the discussion that follows, I will describe the account of perception and phenomenal qualities that Clark endorses as the *non-sensorimotor, skill-based approach*. The approach is framed by Clark as a chapter-long response to the sensorimotor theory rather than a developed account of perception, but in broad terms it advocates a number of positive claims about visual processing and the phenomenal qualities that characterise visual consciousness. The account's claims can be summarised as follows: (i) There are two visual systems, which operate in the way described above; (ii) Ventral stream vision deploys representations that represent properties like type and location rather than sensorimotor law; (iii) The ventral stream deploys these representations in the course of enabling the exercise of skillful capacities like sorting, sifting, classifying, selecting and others, which for shorthand I will refer to just as sorting and sifting; (iv) The properties sorted and sifted are not sensorimotor laws, but types, categories and locations; (v) The appeal to sorting and sifting offers, in principle, a non-mysterious, physicalistic account of the phenomenal qualities associated with conscious perception; (vi) Sorting and sifting are not sensorimotor skills. They are realised by subpersonal systems which are, to a significant degree, informationally encapsulated from the operation of the dorsal stream and motor output systems. This enables them also to be largely insensitive to the fine-grained morphological characteristics of the agent's body.

5.2 - The Policy of Containment: A Conservative Response

Clark (2008) suggests that the sensorimotor approach to phenomenal qualities enters the running alongside competing skill-based accounts that share the explanatory benefit gained by appealing to skills, but lack what he regards as the undue insistence that the skills in question

are necessarily sensorimotor. His response to the sensorimotor account in that chapter is not primarily concerned with phenomenal qualities, however, but with the character of the cognitive processing that realises perceptual consciousness. He argues that the empirical evidence concerning ventral stream perception challenges the sensorimotor account by indicating that perceptual consciousness supervenes on internal states, in particular representations, that stand-in for categories, types and locations rather than sensorimotor contingencies.

To respond, the sensorimotor theory could adopt an approach I will call the Policy of Containment. The response aims to contain Clark's criticisms by restricting the explanatory role played by law-governed sensorimotor interactions to one domain, while asserting that the evidential force of the empirical work Clark cites is restricted to another. Specifically, it contends that the appeal to sensorimotor coupling is not meant to characterise cognition in humans, but solely to account for the phenomenal qualities that feature in perceptual consciousness. As an account of qualia, the sensorimotor approach enters one classic debate, prompted by philosophers like Jackson (1982) and Chalmers (1996), over whether phenomenal qualities are an intelligible part of the physical world. Assuming that phenomenal qualities are physical, it also enters mainstream debates over what sort of physical qualities they are, included in which is the question of which characteristics make one phenomenal quality different from another. On this latter set of questions, the approach enters the field alongside the accounts offered, for example, by Tye (2000), who proposes that phenomenal qualities reduce to representational content; Block (1996), who suggests they reduce to some non-representational neural property; and Pettit (2003), who, in common with sensorimotor theory, offers a skill-based account.

It is important to emphasise that it is not obvious, pretheoretically at least, that an account can address the above topics simultaneous to characterising the cognitive processing that underlies perceptual consciousness. Tye addresses cognition and phenomenal qualities at once, since he proposes that phenomenal qualities are individuated and reduced to the physical by reference to the content of neurally-realised representations, and thereby overtly commits himself to the existence of internal representations in cognition. Clark's account also appears to suggest that there is a one-to-one mapping between the properties represented

subpersonally by the ventral stream and the properties skillfully discriminated in conscious perception. Despite this apparent mapping, however, Clark appeals to quite different things respectively to explain phenomenal qualities and cognition. In the former case, it is the exercise of personal-level skills; in the latter, the activation of representations. In this respect, Clark's account is quite different to Tye's.

On the gloss endorsed by the Policy of Containment, the sensorimotor theory is open to endorsing an account of cognition along the lines of the one advocated by Clark (2008), in particular the claim that visual consciousness is realised by representations that stand in for categories and types. In this respect, the theory endorses a traditional representational approach, although it emphasises that cognitive processing is importantly scaffolded by bodily exploration, a claim which does not conflict with Clark's characterisation of ventral stream vision. To account for phenomenal qualities, however, the sensorimotor theory suggests that we need only focus on regularities in the agent's skill-driven sensorimotor coupling, avoiding altogether the more opaque territory of the brain. On this reading, the theory can be regarded as a kind of *sui generis* hybrid between traditional and radically embodied accounts of mind. The success of this approach depends on showing that there is no incompatibility between its respective accounts of phenomenal qualities and cognitive processing, and, additionally, that there are positive reasons to endorse its account of phenomenal qualities. I address these considerations in turn in the next sections.

5.3 - A Challenge to the Policy of Containment

The dual visual systems account offers one lens through which to understand O'Regan's (2011) claim that conscious perception involves a second-order tracking of the regularities manifest in the perceiver's sensorimotor coupling. We can plausibly suppose that the second-order tracking O'Regan describes is carried out by the ventral stream, which, as Clark contends, fails to play a consistently fluid role in the perceiver's ongoing sensorimotor engagements. The object of the tracking is the bodily coupling that comprises the lowest of O'Regan's hierarchically-organised tiers, and this is plausibly modulated by the dorsal stream. Although Noë (2004) fails to say much about perception at the descriptive level employed by

cognitive science, the conscious sensorimotor understanding featured in his account could itself conceivably be the outcome of the same second-order tracking. The Policy of Containment endorses this picture, but proposes that phenomenal qualities are accounted for by the sensorimotor coupling rather than the second-order tracking.

Notwithstanding its explanatory virtues, this account of phenomenal qualities may fail to persuade if it is not tightly integrated into the sensorimotor account of visual processing. In the next two sections, I examine different versions of the same problem and suggest ways the sensorimotor approach can respond. The drawbacks all follow from the apparent role played by the second-order tracking. As I observed in chapters 3 and 4, the variety of sensorimotor mastery deployed in conscious perception - either Sensorimotor Tracking or Sensorimotor Understanding - is not fully developed by the original statements of the sensorimotor theory. The authors explain what the mastery is mastery *of*, namely SMCs, and claim that it is practical and not propositional (O&N; Noë, 2004). Noë (2004; 2012) also discusses its phenomenology at length. However, neither Noë nor O'Regan offer a clear proposal about how this sensorimotor mastery is realised at the descriptive levels employed by cognitive science or neuroscience. The absence of an explicit proposal about how sensorimotor mastery is physically realised is problematic. Firstly, it entails, straightforwardly, an explanatory limitation: the sensorimotor approach, on the present understanding, comes with no clear naturalistic account of what it is for perceivers to be intentionally-directed toward the objects of conscious perception. Moreover, it invites the view that the sensorimotor account must appeal to internal representation, a reading endorsed by commentators such as Roberts (2010) and Hutto and Myin (2012). From this, all sorts of question follow about the nature of the representations and the explanatory work that they properly do.

5.3.1 - The WAY problem

The suggestion that perception involves representation invites a broad threat to the Policy of Containment, namely the charge that perceptual consciousness, in particular the phenomenal qualities that characterise it, must be accounted for not by the law-like patterns of dependence

that govern the perceiver's sensorimotor engagements, but by properties of the representations themselves. In dialogue with O'Regan, Block makes the following point:

In having a cognitive appreciation of a law involving inputs and outputs, one has to think of or represent those inputs and outputs in some WAY. A machine or a creature from outer space might be able to think of human inputs and outputs in WAYS that do not involve any conscious experience. Alternatively, the WAY might itself be phenomenal—say if our cognitive appreciation is coded in imagery. Given that cognition cannot grasp anything without grasping it in some WAY, the appeal to COGNIZING in explaining sensory qualities smuggles in the very notion that is supposed to be explained. (O'Regan and Block, 2012)

This is a fairly broad point, from which we might identify various specific problems. I will discuss two of them in the remainder of this passage, and a third in the section that follows.

One of Block's complaints is that there are ways of representing sensorimotor contingencies that do not involve consciousness, meaning that an appeal to nothing more than the representation of sensorimotor contingencies cannot be sufficient to account for the existence of consciousness. This is a limitation that the sensorimotor theory can and should accept, given that certain other conditions must be in place for an agent to be conscious at all. O'Regan (2011) asserts that consciousness requires the functional realisation of a faculty described as the *self*, while Noë (2004) expresses sympathy toward the view that consciousness is connected somehow to life; and there may be distinct proposals about consciousness that are equally compatible with the sensorimotor approach. While offering some such background conditions for perceptual awareness is no small task, it is not obvious that an account of perception and the phenomenal character of perceptual consciousness is committed to offering a fully-developed answer to this broader question. With the assumption that appropriate background conditions are in place, the sensorimotor theorist can claim not just that sensorimotor contingencies explain differences and similarities in phenomenal quality, but, more fundamentally, that it is by virtue of representations of sensorimotor contingencies that there are, in general, such things as phenomenal qualities at all.

More seriously, Block's objection suggests that once you have accounted for the presence of conscious qualities by appeal to the representation of some feature, the

explanatory weight is borne by some property of the representation other than what it represents - making any proposal about the particular properties represented a dispensable feature of the account. To assume this suggestion is correct without further argument would be a prejudice. Developing the criticism, if it is not the representation's extension that does the explanatory work, it might be an intrinsic property of the neural vehicle that could inhere even if that vehicle lacked other properties, in particular extrinsic ones, by virtue of which it bears representational content.

Block pursues this line of argument elsewhere in an argument against representationalism about phenomenal qualities. In a thought experiment similar to Davidson's Swampman (1987), Block (1998) imagines that an atom-for-atom duplicate of himself appears, by pure chance, which lacks the causal history needed to possess content-bearing inner states. The duplicate nonetheless undergoes the same phenomenal states, suggesting it cannot be content that yields the phenomenal state. Transposing this objection to the sensorimotor theory, Block imagines a neural duplicate of himself that lacks the relevant ongoing or historical sensorimotor interactions but enjoys identical phenomenal states (O'Regan and Block, 2012). Assuming that this is possible, the scenario appears to suggest that ongoing or historical sensorimotor interactions themselves cannot explain phenomenal qualities, because the neural duplicate does not have them. Similarly, the neural duplicate's brain cannot track or represent the appropriate sensorimotor laws, again because there are no sensorimotor interactions to track or represent.

The sensorimotor theory, on a dynamical systems approach, could respond that duplicating the necessary neural states depends on duplicating the appropriate sensorimotor interactions, as we saw in chapter 3. However, the present aim is to defend a more conservative version of the sensorimotor theory that cannot make use of this argument. Another possible response is to endorse disjunctivism, the claim that veridical perceptual experiences are a different kind of state to non-veridical experiences like hallucinations, even if they are phenomenologically indistinguishable. The disjunctivist response would claim that since the duplicates' states differ in kind, with one undergoing a perceptual state and the other something else, we are entitled to provide a distinct explanation for each. This response may fail to convince, however, on the ground that the sensorimotor account of phenomenal

qualities is not meant to explain perception itself, but the phenomenal character common to both the perceptual and the non-perceptual states.

It is perhaps for this reason that O'Regan endorses a different kind of response. He replies that the neural duplicates share a pertinent feature, but, contra Block, that this is not some set of intrinsic neural properties. The pertinent feature is, instead, the fact that the duplicates' brains are each in the state they would ordinarily be in *if* they possessed the relevant ongoing or historical sensorimotor interactions (O'Regan and Block, 2012). This line of response is insufficient to defend the radical claim that the vehicle of perceptual consciousness extends into the outside environment - notably, it is less robust than the argument offered by Hurley (1998) for ECM, which I discussed in section 3.6. But O'Regan's argument is itself sufficient to defend the claim that a proper explanation of conscious qualities must make reference to features of the outside environment.

5.3.2 - Sense and Sensorimotor Contingencies

On a more interesting reading of the WAY objection, phenomenal qualities are not explained by the extension of the representations, but their intensional content, which the empirical work on the ventral stream perception perhaps suggests is likely not to represent the relevant properties *as* laws of sensorimotor contingency. In this section, with special attention to Tye (2000), I will respond to this worry.

The clearest way to explain the distinction between intension and extension is by reference to the appearance it makes in theories of linguistic meaning. In the standard Fregean example, the terms 'morning star' and 'evening star' share an extension, because they both pick out Venus. However, when we first learn that they pick out the same thing, we learn something new. This is explained by the fact that they differ in intension, an attribute of meaning sometimes described as a mode of presentation. A term's intension is often understood to be a set of truth conditions, although the intension could also be some other feature that meets the requirement of explaining how statements of identity between co-referential terms can be informative. The distinction between intension and extension can also be applied to non-linguistic forms of representation, including internal representation at

the explanatory level employed by cognitive science.

An opponent of the sensorimotor theory might concede that for every difference in phenomenal quality there is a corresponding difference in the pattern of sensorimotor dependence governing your bodily interactions with the object perceived. In so far as this holds true, they might say, phenomenal qualities can be individuated by reference to sensorimotor laws. Nonetheless, the objection continues, it is merely incidental that differences in look are accompanied by differences in sensorimotor dependence. The feature that truly explains a tomato's looking red rather than blue, for instance, is not this unextraordinary feature of the extension, but the intension under which the tomato is represented. The intension may not specify that you stand in a particular relation of sensorimotor dependence to the tomato, but instead that you are looking at an object with a particular set of intrinsic properties, or that you are looking at some object which falls under the descriptor 'red'. To secure an absolute victory over the sensorimotor account, an endorser of this line of response must show that the ventral stream, for instance, represents objects under intensions that determinately do not specify patterns of sensorimotor dependence.

The force of this objection will depend on further technical details about the representational content featured in our account of visual consciousness. As it happens, there is at least one leading representationalist account of phenomenal quality pre-existing the sensorimotor account under which intension poses no threat to sensorimotor theory. Although the account offered by Tye (2000) is superficially quite different from the sensorimotor theory, the variety of representational content he endorses makes a good fit, as it will emerge, with the account of phenomenal qualities offered by the particular reading of sensorimotor theory I am pursuing for the time being. In the remainder of this section, I will focus on Tye's account in order to explain how it defuses the threat posed by intension to the sensorimotor account.

The accounts offered by Clark (2008) and Tye resemble each other inasmuch as they both endorse an important role for internal representation in accounting for cognition and phenomenal character. Tye, however, is not directly interested in the details of human cognition, but in using representational content to account for the latter. Where Clark argues that representations enable the exercise of epistemic skills, which themselves account for

phenomenal character, Tye accounts for phenomenal character directly by appeal to the content of the representations. In the apparent converse of Clark's position, Tye ascribes an enabling role to epistemic skills, suggesting that to feature in perceptual consciousness, the representational content must be poised to impact on beliefs and desires. These positions appear to come apart quite considerably, although I will argue later that Clark's account lapses into something close to Tye's than appearances at first suggest.

Tye's account of representational content is in some ways straightforward. His account of content appeals to no more than a simple co-variance relation between aspects of the brain and environment. He suggests that a neural structure bears content about some feature of the environment just in case it is ordinarily present when and only when that feature of the environment is present. Content in Tye's account is individuated externalistically. To understand this point, recall from earlier Putnam's (1975) classic distinction between internalism and externalism about content. The former denotes that content is fixed entirely by features inside the subject's skin. The latter denotes that it is fixed by properties of the extension. To identify which account is correct, we use a thought experiment featuring physically identical inhabitants of superficially identical planets, Earth and Twin Earth. On Earth, the clear liquid that comes up of taps, fills lakes, and so forth, is H²O. Twin Earth is identical in every respect, except the substance bearing the same superficial features is a different chemical, XYZ. If physical duplicates living respectively on Earth and Twin Earth mean different things when they say 'water', then externalism is true; if they mean the same thing, then internalism is true. And the same distinction can be applied to subpersonal content. On Tye's account, then, a neural state represents water if it is ordinarily present when and only when water is present, and it represents H²O rather than XYZ on account of the properties that water happens to actually have.

Suppose that you see a red object, and your brain represents it as falling under the designator 'red'. Further suppose that being confronted by a red object is the same as standing in a bodily relation to the environment where particular patterns of sensorimotor dependence obtain. If externalism is true, then it may follow that the representational content includes this pattern of sensorimotor dependence. However, this suggestion will not by itself serve well enough as a defence of the sensorimotor theory. Even if intension is determined by properties

of the extension, externalism in itself does not require that the intension is determined by every property of the extension. Tye is motivated to endorse intensional content in the first place in part by the need to account for the fact that we may be aware in perceptual experience of some properties of the extension, but not all. In his example, we might see a vegetable as purple, but not see it as poisonous, even in a possible world where all purple things are poisonous. Intension is needed to account for this, since if all purple things were poisonous, the expressions ‘purple object’ and ‘poisonous object’ would be coextensive.

To see how the sensorimotor account can be defended, we need to look further at Tye’s account. Tye endorses a distinction between *phenomenal* and *epistemic* uses of the word ‘looks’. He argues that the distinction is in evidence in everyday discourse about perceptual states when we talk, at times, about an object ‘looking F to S’ and, at other times, about an object ‘looking as if it is an F to S’. The former locution describes the phenomenal character of a person’s perceptual awareness, while the latter describes the person’s epistemic state. On this account, to say something like ‘the shape looks as if it is a trapezoid to Jim’ implies that Jim has the concept TRAPEZOID, and believes the object he is looking at falls under that concept. Tye implies that this is neither necessary nor sufficient for us to truly say ‘the shape looks trapezoid to Jim’, a locution which implies, distinctly, that a trapezoid shape is part of the phenomenal character of Jim’s perceptual consciousness. Endorsing this distinction allows Tye to suggest that when we talk about looks we are at least sometimes talking about phenomenal character and *not* epistemic situation. On Tye’s account, the shape may look, phenomenally, trapezoid to Jim even if Jim does not possess the concept TRAPEZOID, because the phenomenal look is accounted for by the *non-conceptual* content of his perceptual representations. Tye’s motivation for focusing on what we *say*, or are entitled to say, about the way things look to some third person is, presumably, intended as a way to (1) make it plausible that there is such a thing as a phenomenal look which is not to be glossed in terms of conceptual content and (2) discuss the alleged non-conceptual content of perceptual consciousness without attributing the concepts in question to the subject.

Tye describes epistemic ‘looks’ ascriptions as requiring *hyperintensionality*. Where regular intension allows coextensional expressions to have different contents, hyperintension allows *cointensional* expressions to have different contents. Hyperintension is needed on the

ground that two terms have the same intension so long as they have the same extension by metaphysical or logical necessity. Thus 'H₂O' and 'water' share an intension, as do '55²' and '3025'. These cases are different to the 'morning star' and 'evening star' case, since the shared extensions must be shared in every possible world, while there is a possible world in which 'morning star' and 'evening star' have the intensions they do on this world but are not coextensional, entailing that you are finding out something particular about our world when you learn that they share an identity. The same applies to 'purple object' and 'poisonous object' in Tye's possible world. Nothing turns on whether or not cases where hyperintension is allegedly required cannot, in fact, more accurately be described as instances requiring a variety of plain old intension. It will suffice if we agree that the supposed hyperintension-involving cases are, at least, a distinctive subset of cases involving intension, and I will therefore describe them as involving 'hyperintension' to designate that they belong at a minimum to this special subset, and use 'intension' to describe cases that do not fall under this subset.

Tye argues that the representational content that accounts for phenomenal character is nonconceptual. This reflects the distinction about what we mean in epistemic and phenomenal 'looks' contexts. Crediting Dretske for the example, Tye offers a case in which you see a man in a police uniform, but fail to recognise that he is a policeman. Nothing is strange, Tye observes, about the contention that the man could nonetheless be truly said to look phenomenally to you the way a policeman looks. This lends apparent support to the claim that epistemic and phenomenal looks attributions come apart, and suggests that in the phenomenal case, you need not possess the concept POLICEMAN. If it is indeed appropriate to explain phenomenal look by reference to representational content, then the content in question must therefore be nonconceptual.

Importantly, Tye takes it that the nonconceptual content of perceptual experience cannot be hyperintensional. His argument, which he puts briefly, is that this follows from the fact that hyperintension depends on having "conceptual modes of presentation" (2000, p. 57) which are lacking in the case of nonconceptual content. It is not straightforwardly evident from Tye's discussion why hyperintensional content must be conceptual, but the milder kind of intension that Tye holds is characteristic of nonconceptual content need not be. However,

the view is compelling in light of the following thought, which is perhaps what he has in mind. In the scenario mentioned above, PURPLE OBJECT and POISONOUS OBJECT are co-extensive, because on Earth it just so happens that all purple objects are poisonous, and all poisonous objects are purple. Since the concepts are co-extensive, we cannot identify any physical differences in the way a subject responds to a poisonous object and the way they respond to a purple object. However, we can conceive of a counterfactual circumstance in which the concepts are not co-extensive, and hypothesise that the subject's brain would continue to respond in the relevant way to purple objects that were not poisonous, while failing to respond in the relevant way to poisonous objects that were not purple. The intension is then fixed by the aspect of the extension that the subject's brain reliably responds to, not just now, but in other physically possible circumstances. We cannot, however, make the same move for cases in which the identity is logically or metaphysically necessary, as with water and H₂O, since it is not coherent to hypothesise about how a subject would respond to H₂O that was not water, a scenario that is impossible and perhaps even inconceivable. In the absence of any other readily apparent options for locating hyperintensional nonconceptual content, the only means available to master hyperintensional 'H₂O' and 'water' modes of presentation appears to be by using concepts.

Tye notes therefore that on his account, an object's looking phenomenally red is identical to its looking like an object characterised by the physical properties which constitute being red; as Tye supposes, looking "disposed to reflect such-and-such percentages of light of so-and-so wavelengths" (2000, p. 55). This makes Tye's representationalism quite compatible with the sensorimotor theory's present gloss. On this view, we can take it that second-order tracking of sensorimotor laws involves the activation of representations of the sort described by Tye. In supplement, we merely add that the physical properties represented are not dispositions to reflect such-and-such percentages of light, but patterns of sensorimotor contingency. Importantly, it makes no difference that the sensorimotor laws represented may be more coarse-grained than those discriminated in dorsal stream vision. The sensorimotor account of phenomenal qualities has several important virtues in comparison to Tye's suggestion about the physical properties that underlie colour. Being 'disposed to reflect such-and-such percentages of light' implies that the object activates this disposition when

certain conditions external to the object itself obtain. As such, the disposition is a relational rather than intrinsic property. But Tye's suggestion, which he would of course point out is not intended as a developed proposal, gives no indication of what the other relata is. Sensorimotor theory does offer a proposal on this score, suggesting that the relevant relations are those in which the object stands with ambient lighting conditions and with the agent's body. As a consequence, sensorimotor theory is able to give a better-developed and therefore more principled account of why particular physical states are associated with particular phenomenal qualities, including, crucially, the degrees of similarity and difference between phenomenal qualities. On this basis, there is good reason for an endorser of Tye-style representationalism to adopt the sensorimotor account of phenomenal qualities.

On the present view, the sensorimotor theory emerges as a hybrid of skill-based and representationalist approaches to qualia. The above discussion has shown we can claim, without contradiction, that sensorimotor skill does the main explanatory work in accounting for conscious feel, even though representation is a necessary condition, along with other enabling conditions, for being in states characterised by the relevant phenomenal qualities. On this proposal, the ventral stream plausibly groups together fine-grained patterns of sensorimotor dependence under representations of coarser-grained categories and types, which it deploys for the purposes both Clark and O'Regan recommend. Since the representations lack the strong kind of intension required to represent the relevant features, determinately, *as* categories and types and *not* patterns of sensorimotor dependence, the option remains to individuate the associated phenomenal qualities by reference to sensorimotor law. I have shown in brief that the sensorimotor account, on the present understanding, brings explanatory benefits as an addendum to Tye's representationalism, and another such benefit will become evident in the next section when I discuss the Inverted Earth thought experiment. Returning to the central focus of this chapter, it remains necessary to show that the sensorimotor account is not just coherent but actively compelling when considered alongside the alternative, non-sensorimotor, skill-based approach.

5.4 - Mary and Inverted Earth: A Sensorimotor Response

The above discussion has illustrated that we could quite consistently hold that perceptual experience is realised by the deployment of representations which themselves are not intimately involved with bodily movement, while also holding that the phenomenal character of the experience so-realised is best explained by the sensorimotor laws that govern the perceiver's bodily interactions. Nonetheless, it may seem more natural to suppose that phenomenal character is accounted for by the personal-level skills the representations enable - as Clark (2008) argues, the sorting and sifting of categories and types for the purpose of judgement and action-planning. In this section I offer independent reasons for endorsing the sensorimotor account of phenomenal qualities. Importantly, there is a range of conceptual demands an account of phenomenal qualities must meet that turn on considerations superficially distinct from the character of the representations we are supposing enable perceptual consciousness; and it will become evident as I proceed that sensorimotor and non-sensorimotor skill-based accounts are not equivalent in their ability to meet these conceptual demands. Attention to some classic thought experiments will help show that there is reason to think the sensorimotor approach to phenomenal qualities is more effective at meeting these demands than its competitors.

Consider first the debate over the 'inverted spectrum' thought experiment. Generically, inverted spectra thought experiments feature two perceivers, one whose colour phenomenology is inverted relative to the other, but who exhibits an identical repertoire of behaviour. Block (1990b) has used a version of the thought experiment in an attempt to refute the view, endorsed by Tye (2000) and many others (e.g., Lycan, 1987; Dretske, 1995), that phenomenal qualities reduce to representational content. Block and his opponents share the background assumptions that phenomenal qualities are physical, and that there is no deep-seated conceptual obstacle to accounting for representational content in physicalistic language either. Block imagines a planet which is identical in all respects to our own apart from the fact that the colours of objects are complementary to the colour they possess on Earth, meaning the sky is orange, the grass red, and so forth. The inhabitants of Inverted Earth correctly describe the sky as 'blue', and the grass as 'green', since the meanings of those words

on Inverted Earth differ appropriately from their meanings on Earth. Block asks us to imagine that we put on a pair of colour-inverting glasses and travel to Inverted Earth. The inverting effect of the glasses means that things appear phenomenally to us on Inverted Earth exactly as they do on Earth without the glasses. Yet, at least on externalist accounts of representational content, the content our experiences possess must now differ in line with the differing properties of the objects we are now in causal contact with. This appears to show that phenomenal qualities cannot be adequately accounted for by representational content, at least so long as the content is individuated externalistically.

The sensorimotor approach (on the present gloss) agrees with the representationalists that phenomenal qualities are necessarily individuated by properties represented when we perceive. However, sensorimotor theory avoids the threat that the Inverted Earth case poses to this view, since it is not committed to the suggestion that the intrinsic properties of the objects we come into perceptual contact with could not vary while phenomenal character stays the same. The approach instead holds that phenomenal character will stay the same between Earth and Inverted Earth so long as the relevant patterns of dependence between motor output and sense input are preserved. If this is correct, the inverting goggles could only do the job they are supposed to do if they somehow preserve the patterns of output-input dependence that obtain on Earth. In this event, the phenomenal character of your experience on Inverted Earth is explained, amongst other things, by properties of the objects perceived and properties of the goggles (although, as Myin, 2001, points out, inverting goggles that preserve the relevant patterns of sensorimotor dependence may be a physical impossibility).

Suppose that objects with intrinsic property X are, in every possible world, co-extensional with objects with which a perceiver embodied in a particular way stands in sensorimotor relation S . This means that a neural representation, lacking hyperintension, does not represent only X or only S , but represents both simultaneously. On Inverted Earth, the equivalent object has intrinsic property Y instead of X . Because you are wearing appropriate goggles, neural implants, etc., you stand in sensorimotor relation S_2 to the object. S_2 is, in one respect, non-identical to S , because it is realised in part by an object with different intrinsic properties. However, S and S_2 correspond isomorphically in a way that X and Y do not, entailing that there is a level of description at which S and S_2 are identical. On Inverted Earth,

there is no way to specify determinately whether the neural representation bears content about Y or S2. However, by appealing to the resemblance between S and S2, we can explain why the phenomenology is preserved in spite of the difference in representational content. This response gives a representationalist like Tye further reason to incorporate the insights provided by the sensorimotor account.

The non-sensorimotor skill-based approach offers a different kind of response to the Inverted Earth case. The approach has no need in the first place to insist that phenomenal qualities are individuated by properties of the objects we perceive. All that matters is that, for whatever reason - whether because of the goggles, or some property of the environment - we are able to make a suitable range of discriminations, which do not merely involve judgements that one property is nonidentical to another, but graded judgements about similarity and difference along various scales, for example hue, saturation and brightness. So far, both skill-based accounts offer in principle an equal improvement on representationalism about qualia, although I have established so far that they are not equivalent, since the sensorimotor account appeals, in part, to properties of the outside environment, while the non-sensorimotor skill-based account does not obviously do so.

Consider now a debate concerning the more general question of whether phenomenal qualities are physical at all, and the vivid challenge posed to physicalism by Jackson's Mary's room thought experiment. To recap, the story features Mary, a vision researcher, who is locked in an entirely black and white room and never comes into direct visual contact with any coloured objects. Via black and white TV screens, Mary studies vision science and comes to know every fact that there is to be known about the physical world. Yet when she is finally released from the room, and shown a red rose, Mary, it is alleged, learns a new fact; namely, what it is like to see something red. Since Mary, by hypothesis, already knows every physical fact, Jackson concludes that she learns a non-physical fact, and hence that physicalism about qualia is false (Jackson, 1982). The sensorimotor theory does not come with an entirely prefabricated response to the Mary scenario, but it does invite certain lines of response, and precludes others. The theory says that the look of red can be characterised by reference to a bodily relation the perceiver stands in to the world, which is itself a physical relation, meaning the theory is committed to arguing somehow that phenomenal qualities are physical.

More importantly, the theory is independently committed to the view that to stand in this relation is to possess and deploy a certain kind of skill or practical knowledge. With this in view, it is natural for it to endorse the suggestion (see, Lewis, 1983) that to ‘know what it is like to see red’ is to possess a kind of practical rather than propositional knowledge. Mary’s new knowledge, in this light, plausibly consists of the ability to deploy, at will, the same sensorimotor mastery she deploys to perceive a red object even on occasions where there is no red object present. This is an act of the imagination, and involves naturally extending the sensorimotor theory to encompass imagination in addition to perception. The sensorimotor account is able to make better use of the appeal to practical knowledge to account for the Mary case than accounts appealing to practical knowledge in a more *ad hoc* way, because the theory makes use of practical knowledge to account for perception anyway.

This response gives the sensorimotor account an advantage over the non-sensorimotor skill-based account. In line with Clark’s proposal, that account is committed to the suggestion that seeing something phenomenally as red consists in nothing more or less than having the ability to judge that it is red, or plan one’s actions with sensitivity to the fact that it is red. Mary, locked in the room, possesses all the propositional knowledge about physical facts there is to be possessed. If deploying this propositional knowledge in thought depends on possessing practical knowledge, as Ryle (1949) contended, there is no reason to think she lacks this ability. With this in mind, she can discriminate in judgement the same range of properties before leaving the room as she can after leaving it. If the epistemic skill theorist thinks it is the capacity to make propositionally-structured judgements that makes the difference, they are committed to the view that Mary knows what it is like to see red before leaving the room, a position endorsed by Dennett (1991), but which would, at best, add a substantial and perhaps unpalatable additional commitment to the epistemic skill-based approach.

The non-sensorimotor skill theorist could instead suggest that the capacities involved in ventral stream perception involve neither the deployment of propositional knowledge nor the exercise of fine-grained sensorimotor skill, but a different kind of practical skill. Ward (2009), for instance, suggests that the skill in question is the planning of agential action. If this amounts only to an appeal to coarse-grained rather than fine-grained sensorimotor skill, this is

the position I argued earlier is not, as Clark suggests, incompatible with the sensorimotor theory. Alternatively, the skill might be a capacity for something other than bodily movement or conceptual thought. It is not self-evident what we should take this capacity to be. The most obvious possibility is that it is a capacity for nonconceptual thought. In line with the understanding of nonconceptual content endorsed, for instance, by Tye (2000), a thought of this kind is a conscious state featuring content only characterisable by concepts that the subject does not possess.

However, the appeal to personal-level nonconceptual content is baseless unless some independent considerations make it plausible and intelligible that we actually undergo states of this kind. We cannot make the existence of such states intelligible by appeal to visual phenomenology, as this would be circular: we are trying to make the visual phenomenology intelligible by appeal to the content. The only remaining option is to appeal to the subpersonal content borne by representations realised by the ventral stream or another part of the brain. To explain what it is for a person to experience one patch of red as darker than another, for instance, the present approach would have to say that the person's ventral stream represents it as darker. Notice that the epistemic skill-based account amounts, on this proposal, to nothing more than representationalism in a different guise. Lacking the modification that sensorimotor theory makes on its present representationalist construal, this threatens a return to the Inverted Earth problem, namely the possibility that the discriminations you make could vary while your experience's phenomenal character remains the same.

5.5 - Relation to Rupert's 'Massively Representational Mind'

On its present gloss, the sensorimotor account endorses something similar to Rupert's (2011) notion of *internal-external content*. Typically, representational stories pitched at a subpersonal explanatory level propose that the brain traffics in representations which stand-in purely for outer features of the subject's environment. This holds true in the account offered by Tye (2000), and could correctly characterise Clark's (2008) claim that the ventral stream represents types, categories and locations. Rupert's relatively novel suggestion is that parts of

the brain frequently represent other parts of the brain, and parts of the extra-neural body. In one sense, this content is internalistically individuated, since it is individuated only by features inside the skin of the subject. In a more interesting sense, however, the content is externalistically individuated, because it is fixed entirely by properties of the extension and not by properties internal to the neural state that does the representing. As presently construed, the sensorimotor account suggests something similar. The contents borne by representations in the ventral stream are fixed entirely by properties of the extension. The extension, in this case, is not an external feature but a pattern of dependence governing the perceiver's bodily interaction with the environment. This is not, as with Rupert's internal-externalist content, a purely inner state, but is not wholly external either. Instead it is a relational property that holds between the perceiver and their environment. As I explained in the last section, one benefit of adopting this claim is its ability to overcome the Inverted Earth problem.

It is worth looking in more detail at Rupert's model, the 'massively representational mind', and considering its implications for the sensorimotor account. Rupert offers a very austere account of representation and content, suggesting that X represents Y in the event that it reliably co-occurs with Y. X here is some neural state, and Y is a feature of the outside environment, body or brain. Where the thing represented is another part of the brain, it may be the neural state itself, or it may be the content carried by that neural state. Content is in every case individuated solely by the properties of its extension. Moreover, bearing content is sufficient on Rupert's model for constituting a representation. Rupert makes no demand that the contents fulfil, for example, any particular functional role. In offering such a liberal account of 'representation', Rupert finds that there are representations everywhere; the brain is massively representational. One important feature of the model is that the brain employs a great deal of 'redundant' content, which is content borne at once by multiple representations. He addresses the dual visual systems to illustrate this feature, arguing that the systems carry some content in common, albeit using different representational formats and in the service of performing distinct functions. Rupert does not claim that the ventral stream bears internal-externalist content; he merely suggests that the ventral and dorsal streams each represent the same features of the subject's external environment. However, it is worth considering the dual visual systems in light of Rupert's internal-externalist content and some

of Clark's claims.

One observation Clark makes is that the ventral stream is indeed sensitive to efferent copy from motor output regions in addition to sense input, suggesting it is sensitive to the relation *between* motor output and sense input, i.e. the patterns of dependence appealed to by the sensorimotor account. His argument, however, is that the ventral stream is sensitive to motor output states only during learning phases, when it is acquiring new information about the significance of various kinds of sense input. In the normal course of perception, however, he claims that the ventral stream is sensitive only to sense inputs. The representations it constructs in the course of ordinary perception are, meanwhile, insensitive to *fine-grained* patterns of sensorimotor dependence. Clark's opposition here depends on what he describes as "the best *functional and architectural explanation*" (2008, p. 192) for the available evidence. This is important. The fact that the ventral stream is insensitive to fine-grained differences in sensorimotor law does not mean that it is insensitive to sensorimotor laws. Moreover, the distinction between learning and ongoing sensitivity is only important if we focus on neural architecture. Focusing only on content, where we take this to be explained by co-variance, there is no theoretical difference between a case where the ventral stream co-varies with particular sensorimotor laws because it is responding in part to efferent copy from the motor areas, and a case where it co-varies with sensorimotor laws even though it is responding directly to sense inputs alone.

We could, on Rupert's model, propose that the ventral stream groups the fine-grained patterns of sensorimotor dependence governing dorsal stream activity into coarser-grained representations corresponding to types, categories and relative locations - thereby altering the representational format - and puts them to use for reasoning and action planning. Despite this difference in format and function, the ventral stream represents or borrows content from the dorsal stream. This proposal therefore assigns to the dorsal stream an explanatory role in individuating the content borne by ventral stream representations, even if it is ventral stream representations that serve as the primary vehicles of visual consciousness. If the content by itself explains phenomenal character, then activity modulated by the dorsal stream plays a relevant role in accounting for phenomenal character. And as I argued in the previous section, Clark's preferred variety of skill-based account in the end depends crucially on the content

borne by the ventral stream representations to explain which properties have been skillfully discriminated.

Simultaneous to co-varying systematically with the operations of the dorsal stream, the ventral stream also co-varies with categories and types present in the outside environment. This may lead an endorser of Clark's view to repeat the claim that the ventral stream's co-varying with features of the dorsal stream is merely a contingent feature, while its directly co-varying with features of the subject's outside environment is the pertinent co-variance relation when specifying the content borne by ventral stream representations, and so accounting for the acts of discriminations that explain phenomenal character. In this event, sensorimotor contingencies play no essential role. But nothing in co-variance based accounts of content, like Rupert's, shows that the content of ventral stream representations are not sensorimotor contingencies. The claim that sensorimotor contingencies are the contents is certainly not arbitrary or unmotivated, either, considering that, as Clark appears to allow, the dorsal stream is appropriately described as enabling the perceiver to exercise skillful mastery of sensorimotor laws, and further that the ventral stream, at least during learning phases, does not just happen to co-vary with these patterns but is directly responsive to both to sense input and efferent copy.

It is therefore possible, at least, to characterise ventral stream content as patterns of sensorimotor contingency. There is at least some positive reason to do so in light of the need to offer a coherent account of phenomenal qualities, which may appeal to different features than any accompanying account of cognition. The ability of the sensorimotor model to respond adequately to the Inverted Earth problem makes the sensorimotor perspective in one respect more coherent. On Inverted Earth, the ventral stream begins to reliably co-vary with, therefore represent, different features of the external environment, and therefore enables a different set of epistemic discriminations, where these discriminations are held to involve discriminating types and categories in the external environment. Yet phenomenal character remains the same. However, the ventral stream continues to co-vary in the same way with dorsal stream activity. This makes dorsal stream activity a better candidate for the content borne by the ventral stream than the outside environment, at least when it comes to explaining phenomenal qualities.

5.6 - Chapter Conclusion

This chapter has observed that phenomenal qualities and cognition are separate topics, comprising different sets of question and having the potential to demand distinct, or at least partly distinct, lines of explanation. I have argued that the sensorimotor account of phenomenal quality is not obviously opposed to a representationalist story about cognition, and that the account comes with certain explanatory benefits when contrasted with other available accounts of qualia. I further argued that the epistemic skill-based account of phenomenal qualities, endorsed by Clark, collapses either into an appeal to sensorimotor skill or into a more or less traditional representationalist story, albeit one lacking the explanatory advantage that comes from claiming the properties so represented are patterns of sensorimotor dependence.

Providing a watertight defence, insofar as this is possible, would require surveying further accounts of representational content, in particular internalist accounts, and exploring in more detail ways they can respond to issues arising in debates over phenomenal qualities, in particular Block's Inverted Earth case, which has been a key feature of my discussion in this chapter. Within the scope of this chapter, I have shown, modestly, that the sensorimotor account is not obviously misguided and has some advantages over Clark's account as an account of phenomenal qualities. In the next section, I offer a more robust defence. I will continue to endorse the suggestion that the ventral stream plausibly tracks and perhaps represents lower-level patterns of sensorimotor dependence. However, I will suggest that conscious qualities are accounted for not by the purported contents carried by processes in the ventral stream, but by the coarser-grained sensorimotor skills that this tracking enables.

6 - Sorting, Sifting and Other Sensorimotor Skills

This chapter offers a more radical response on behalf of the sensorimotor theory to the objection levelled by Clark (2008). In this chapter I will suggest that ventral stream representations, if they exist, and whether or not they stand-in themselves for patterns of sensorimotor dependence, enable perceptual consciousness by enabling the perceiver to exercise varieties of sensorimotor skill, which subsume the epistemic skills appealed to by Clark. It is the exercise of this sensorimotor skill that constitutes perception and accounts for phenomenal character. The force of the defence will be based on the sensorimotor theory's ability to account for the intentional directedness and phenomenal character of visual consciousness. This will involve subsuming a range of different skills under the heading of sensorimotor skill, but rather than being an unwarranted instance of, as Clark puts it, the "enactive framework wagging the empirical dog" (p. 192), I will suggest it is a soundly motivated demand that subpersonal theories not mis-describe the relation between, on the one hand, neural and bodily states, and on the other, the personal-level phenomenal states they realise.

6.1 - Avoiding Zombie Intuitions

To begin, let's consider yet another famous thought experiment. Chalmers (1996) is an opponent of the sensorimotor and all other physicalistic approaches to phenomenal consciousness, since he argues that phenomenal consciousness, in general, cannot be physical. To motivate this claim, he argues that we can conceive of a world physically identical to our own and containing creatures that are physically identical to ourselves but entirely devoid of phenomenal consciousness. He call these creatures zombies. On the ground that we can conceive of such a world, Chalmers claims that such a world is metaphysically possible. And if the relevant physical states and processes could exist without the relevant phenomenal states, then the phenomenal states cannot be identical to the physical states. We could attempt to resist Chalmers' argument from the outset by denying that all conceivable situations are possible, or by denying that zombies are even conceivable. Whichever strategy we pick, it

behooves the physicalist so far as possible to reply not just with negative claims but also with a positive proposal about the sorts of physical states that do constitute phenomenal consciousness. Clark (2008) is explicit that he approves of skill-based accounts because they offer a conception of qualia in which we cannot conceive of a possible world of the sort Chalmers describes: on a skill-based account, there is simply nothing more to undergoing a particular phenomenal state than exercising the relevant skill.

Accounts of cognition, some of which comprise explanations of phenomenal character, frequently advocate that certain relationships hold between entities at different explanatory levels. For instance, they might say that representational content, at a subpersonal explanatory level, supervenes on particular kinds of neural state. They may also, like Clark, say that visual consciousness supervenes on the deployment of particular kinds of representation. A supervenience relation implies that the thing doing the supervening cannot vary without the supervenience base varying. At times, a supervenience relation might be metaphysically necessary, meaning it holds in every possible world, or nomologically necessary, meaning it cannot fail to hold in this world. However, there is no general obligation for a supervenience relation to imply anything so strong. For instance, it would be informative and plausibly true to claim that visual consciousness supervenes on ventral stream activity in presently living humans, on the ground that visual consciousness in humans cannot vary without ventral stream activity varying. However, it would be unjustified to claim, with confidence, that future humans could not in principle use technological artifacts or a different brain areas to realise the same kind of consciousness.

Accounts of phenomenal quality are also entitled, for some explanatory purposes, to make supervenience claims without committing to their obtaining in every possible world. To show that zombies are impossible, however, an account of qualia must offer something more robust. At a minimum, it needs to show that a given physical property, or physically-realizable functional property, is metaphysically sufficient for an appropriate phenomenal state to obtain. Logically speaking, metaphysically sufficient conditions of this type are enough to entail that there is no possible world containing zombies. However, metaphysically sufficient conditions for phenomenal consciousness would not be enough to discount the inverse view that there is a possible world where phenomenal consciousness exists without any

attendant physical state. There is also the question of how to convince present day dualists subscribing to the zombie intuition that a given set of conditions is indeed metaphysically sufficient for phenomenal consciousness, considering we cannot travel to other possible worlds to find out. To offer a convincing reply to dualists, we must make it convincing that the physical states under consideration are not just sufficient for, but *identical to*, the appropriate phenomenal states. This will make it inconceivable that possible worlds exist where the physical states in question are not metaphysically sufficient.

One view is that psycho-physical identities require metaphysically necessary as well as sufficient physical states. To identify metaphysically necessary conditions, we must, of course, be indifferent to facts that are merely contingent. Assuming, for example, that the ventral stream itself only plays a contingent role in realising the cognitive processing that underlies visual consciousness, the metaphysically necessary conditions will not include having a ventral stream. An endorser of a Clark-style position would claim, more plausibly, that physically-realizable functions are what count. Thus, human-style visual consciousness depends by metaphysical necessity on representations carrying out the functional roles performed in this world by the ventral stream. These include playing an appropriate functional role in the wider cognitive economy, and bearing content that is systematically insensitive to occurrences at the sensorimotor periphery. If Clark were indeed arguing that these states are metaphysically necessary for visual consciousness of the sort undergone by human perceivers, he would be making a strikingly strong commitment. On this basis, he would have to claim that a subject undergoing the same functionally-defined inner activity but altogether lacking the capacity to integrate this activity with skillful behaviour will still experience all or most of the same visual states we do. He would also have to claim that a subject lacking the functionally-defined inner states we possess, but reliably manifesting an identical repertoire of skillful bodily behaviour, necessarily lacks the kind of visual consciousness we enjoy. These claims might be correct, but they go significantly beyond claims that can be straightforwardly justified by attending to the workings of the ventral stream.

To make the sensorimotor approach a more compelling alternative, we must be careful how we frame it. The approach would be committing itself to something considerably less palatable than the above proposal if it claimed that bodily activity is sufficient for visual

consciousness even in a metaphysically possible world where the perceiver lacks a brain. It would not be compelling, either, to insist that bodily activity or activity in motor areas of the brain are necessary for visual consciousness in every possible world. What the sensorimotor approach ought to say, instead, is that the metaphysically necessary and sufficient condition for having visual consciousness characterised by the appropriate phenomenal qualities is the deployment of inner states that make an intelligible contribution to enabling appropriate bidirectional patterns of systematic dependence between sense inputs and motor outputs. Nothing in Clark's (2008) critique demonstrates that ventral stream activity could not be made intelligible this way. If successful, the present approach falls short of establishing the strong claim that the extra-neural body and environment play a metaphysically necessary role in realising perceptual consciousness. But, crucially, it offers a way to resist Clark's claim that the skills which constitute visual consciousness are not sensorimotor skills. The present picture has the advantage of allowing that perception *could* be sensorimotor, while the Clark-style position must actively deny that a creature lacking the right internal states could undergo the kind of visual consciousness humans do, even if they possess the appropriate sensorimotor skills

If there is a need to choose whether sensorimotor or non-sensorimotor skills are more plausible candidates as metaphysically necessary conditions for perceptual consciousness, the sensorimotor account thence has an advantage. However, Clark could reject the need to make this choice. Instead, he might claim that we should endorse pluralism about the skills that comprise the metaphysically necessary conditions. On this understanding, the skills we must appeal to in order to dissolve the zombie intuition are neither sensorimotor nor non-sensorimotor, but non-specific capacities to make the appropriate discriminations. The pluralistic perspective entails that the generic skills in question could in principle be constituted, in different cases, by quite different capacities. In this light, Clark's objection to the sensorimotor account can be taken as targeting the claim that the skills in question are constituted by sensorimotor skills in this world. This returns the advantage to Clark.

However, certain considerations blunt the force of this response. It is an open question whether sensorimotor and non-sensorimotor skills belong to one broad kind. The capacity to make skillful bodily movements and the capacity to represent types, categories and

locations are, on the face of it, quite different kinds of capacity. The only thing they necessarily have in common is that they involve responses that co-vary systematically with aspects of the outside environment. But lots of things exhibit this kind of sensitivity and do not justify ascriptions of 'skill'. A skill-based account of qualia must therefore appeal to some further criterion. In Clark's account, this is the content and functional role of the representations deployed by the ventral stream, however these notions are themselves glossed. In the sensorimotor theory, the skill is ascribed by properties of the outward behaviour, a claim I will elaborate on in the next chapter. Once these criteria are added, there is no guarantee that the skills appealed to respectively by sensorimotor and non-sensorimotor skill-based accounts have any common factor that would justify our thinking of them as two different ways the same generic skills could be constituted. In offering psycho-physical identities, it will not do, then, to claim that skills are the metaphysically necessary conditions for perceptual consciousness without specifying in more precise terms what sort of skill we have in mind. As I have argued, sensorimotor skills are a more compelling candidate, since they accommodate the possibility that perceivers lacking our functionally defined inner states may nevertheless enjoy similar types of experience.

Adopting a different tack, we might find that psycho-physical identities fall short of demanding metaphysical necessity. We could instead make a case that there is simply nothing more to undergoing a given phenomenal state than exercising a particular skill on the ground that exercising the relevant epistemic skills is, in this world alone, even in one subject alone, constitutive of undergoing the phenomenal state. The constitution claim is less strong than the metaphysical necessity claim, since it allows that a given experience might alternatively be constituted of quite different things. For example, there might be nothing more to your travelling to work today than your cycling from A to B, but tomorrow there might be nothing more to your travelling to work than your driving from A to B. Constitution claims go beyond mere sufficient conditions, however, by showing that there is nothing more to consciousness than exercising the appropriate skills, in this world at least. In this light, an endorser of the non-sensorimotor, skill-based account might accede to what we could describe as a disjunctive account of the skills that are metaphysically necessary for visual consciousness. This resembles the pluralistic account of skill I considered above, but rather than claiming that sensorimotor

and non-sensorimotor skills belong to a single broader class, it simply claims that the relevant skills could either be sensorimotor skills OR non-sensorimotor epistemic skills. This avoids the need to claim that a possible creature possessing sensorimotor skills like ours but lacking the relevant inner representational states could not possibly share human-style visual experience, but allows Clark to maintain that in human perceivers, in this world, the skills are non-sensorimotor.

This view has more promise. However, it depends on the claim that there is a genuine disjunction in this world between sensorimotor skills and the epistemic skills apparently enabled by the ventral stream. To show that such a disjunction obtains, Clark would have to show that there can be a difference in epistemic skill without a difference in sensorimotor skill, and vice versa. To prove the sensorimotor account wrong, Clark would also have to show that phenomenal qualities vary in line with the former and not the latter. Clark (2008) does not establish this, only demonstrating that not every difference in sensorimotor skill comes with a difference in phenomenology, an observation the sensorimotor account can accommodate.

If the sensorimotor perspective is correct, there is a constitutive link between perceptual consciousness and certain bodily capacities. It might prove that the bodily capacities in question are constituted in turn by processing involving that deployment of representations that stand in for types, categories and relative locations. But it would, in this eventuality, be the bodily capacities themselves that rightly overcome the zombie intuition. In the next section, I will defend the conceptual claim that the epistemic activities described by Clark, the sorting and sifting of types and categories, in fact constitute sensorimotor skills.

6.2 - Sorting and Sifting as Sensorimotor Skills

Clark (2008) argues that if conscious perception, as realised by the ventral stream, involves sorting and sifting of types and categories in place of the fine-grained sensorimotor behaviour enabled by the dorsal stream, then the sensorimotor approach is thrown into significant doubt. This only follows, however, if sorting and sifting (and the other 'epistemic' activities mentioned, like classifying, selecting and choosing) are not varieties of sensorimotor activity themselves. If they are, then it is straightforward that they draw on mastery of the laws of

sensorimotor contingency. Clark does not explicitly address this possibility, but his rejoinder to the sensorimotor approach implies that sorting and sifting draw purely on internal processes, in particular the activation of internal representations, and not sensorimotor mastery. In this section, I am going to attempt a defence of the sensorimotor approach by showing that sorting and sifting are themselves plausibly best understood as varieties of agent-level sensorimotor activity.

Sorting and sifting make an interesting case study, since they are paradigmatically cognitive activities, while at the same time being, ordinarily, activities carried out through bodily engagement with the environment. This point dovetails with the main focus of Clark's (2008) book, which argues that many of the cognitive processes traditionally held by cognitive scientists to depend solely on neural activity also depend on the agent's bodily engagements with the outside environment. Clark claims that in some cases, the substrate of cognition actually extends into the outside environment. In other cases, he merely claims that to correctly understand how a subject performs a given task, we must make explanatory reference to the subject's embeddedness within a specific environmental niche.

A reasonably paradigmatic example of the latter, offered by Clark, concerns a technique employed by a practised bar tender (p. 62). Each drink on the menu is served in its own distinctively shaped glass, and as the orders come in, the bar tender places appropriate empty glasses in a spatial sequence corresponding to the order in which the drinks are to be served. This saves the bar tender the trouble of trying to remember which drinks are to be served and in which order. Clark stops short in this case of claiming that the bar tender's manipulation of external props actually constitutes part of the subject's cognitive processing, but claims that the trick contributes in an important way to the bar tender's performance of the task, while altering the character of the cognitive tasks that must be carried out by the subject's brain. It does not matter too much for my present point whether the bar tender's interaction with their niche is supposed to itself be cognitive. It suffices to notice that the subject, here, is performing tasks of the kind allegedly carried out by ventral stream perception - classifying, selecting, choosing, and so forth - by means of their sensorimotor coupling with the outside environment. It is striking that from a commonsense perspective, this is no surprise, since in everyday language, if not the language of cognitive science, activities like

sifting and sorting are almost always activities carried out through the bodily manipulation of external props.

By extension from the bar tender case, it is worth entertaining the possibility that the sorting and sifting of types, categories and locations that Clark thinks underlies perceptual consciousness are also constituted by bodily interaction with the environment. This proposal is not altogether outré by the standard of the dual systems model. The model describes dorsal stream activity, which is not supposed to come with any attendant consciousness, as a kind of 'vision', reserving the term 'perception' for cases associated with conscious awareness. If the dual systems and sensorimotor theories are agreed that there is an unconscious kind of vision best characterised as a sensorimotor process rather than by an orthodox representational model, then relative to accounts that reserve 'vision' for states or activity realised solely by the deployment of internal representations, it is a lesser step to contend that conscious perceptual awareness is also constituted by the exercise of sensorimotor skill.

With the above points in mind, the sorting and sifting identified with perceptual awareness is plausibly constituted - in *some* cases - by successful behaviour directed at a practical goal. For example, your sorting red from green might involve nothing more than your appropriately pulling your car to a stop in front of a red traffic light. It would be useful, however, to give some account of how such behaviour is actually achieved. Such an explanation can be given in part, by reference to bodily activity directed toward an epistemic goal, for example angling your head upward to see the traffic light. This sort of ongoing re-alignment of bodily posture for the purpose of gathering information also plausibly constitutes part of the process of sorting and sifting. Although perceptual consciousness does not usually involve active manipulation, with your hands, of external props, a core analogy with the bar tender case is maintained by virtue of the fact that the perceiver, in line with O'Regan's idea of the world as an outside memory, makes persistent changes to their bodily relation with the environment in order to achieve both epistemic and practical goals. The fact that bodily movement plays an important role in perception does not, of course, prove that the extra-neural activity actually constitutes, rather than merely assists, the exercise of the epistemic skills Clark talks about. However, my proposal ought to make it at least conceivable that identifying colours, shapes, locations and so forth is best characterised by reference to a

bodily process.

This suggestion should not be confused with the implausible thesis that perceptual consciousness is realised by nothing more than outwardly-observable bodily behaviour. The proposal is that activities such as sorting and sifting are constituted by sensorimotor coupling carried out by the whole organism, drawing on both neural and non-neural activity. Where outward activity is present, this alone is sufficient, sometimes, to constitute an act of sorting or sifting, provided it is enabled in some intelligible way by the subject's inner activity. To account for discriminations made in perceptual consciousness that go beyond those which can be individuated by outward bodily activity, including documented experiments where subjects have had their whole bodies, including eyes, completely paralysed while enjoying a restricted degree of perceptual awareness anyway (Whitham et al., 2011) a different story must be given. However, these cases could be accounted for by appeal to the agent's *preparedness* to make skillful bodily movements.

It is conceivable, of course, that in order to be prepared to act, the subject must deploy internal representations, such as forward models or action-oriented representations like the ones endorsed by Ward, Roberts and Clark (2011). This is not necessarily a problem for the sensorimotor theory, which on some formulations accepts that subpersonal representations play a role. In this event, perception can still be understood constitutively as a temporally-extended process, which, aside from exceptional circumstances, incorporates bodily movement, which forward models help enable. This is compatible, for instance, with O'Regan's (1992) original presentation of the outside memory idea, which suggests that bodily exploration is assisted by sparse representations of objects and their egocentrically-specified location in the subject's environment.

However, recent work on preparatory activity in the motor cortex supports the suggestion that neural representations are not necessarily involved in preparing for behaviour. It is agreed that neural activity takes place in the motor cortex prior to movement which helps enable the subject to move their eyes, limbs (etc.) the appropriate distance, direction and speed. A popular view has it that certain neurons become weakly active in preparation for a movement, and that the movement takes place when that neural activation reaches a certain threshold, suggesting that the preparatory activity acts as a representation of the movement

that is about to follow (e.g., Tanji and Evarts, 1976). However, this assumption is challenged by recent empirical work. One study recorded neural activity in monkeys prior to and during specific reaching tasks. The study found that the preparatory activity correlated with a given movement was quite different to the neural activity that occurred while the same movement took place. This does not establish that there is no representation going on, but it counts against the representational view I just described. Moreover, it is expected by a mechanistic account, in which the preparatory activity does not represent the movements that are about to follow, but initiates a dynamical system in which the body movements takes place (Churchland, Cunningham, Kaufman, Ryu and Shenoy, 2010). Given that the ventral stream has the function of enabling coarse-grained sensorimotor behaviour, it is conceivable that its activity could also be understood as non-representational preparatory activity for embodied acts of sorting and sifting. Moreover, this shows that perceivers could be engaging the relevant bodily capacities even where the capacities are not manifest in outwardly observable bodily movements.

Sorting and sifting are, in general, naturally understood as agent-level sensorimotor activities. In the domain of conscious perception, it is tenable, at least in some cases, that sorting and sifting are constituted by agent-level sensorimotor coupling. This ought to defuse the suggestion that where perceptual processing involves sorting and sifting, it does not involve sensorimotor skill. My discussion has aimed to offer some positive motivation for the view that sorting and sifting in perception is indeed always a sensorimotor process. It remains for the non-sensorimotor skill-based approach to show that sorting, sifting, classifying and selecting can be explained purely by reference to what the brain does without committing a category mistake.

6.3 - Homunculi that Sort and Sift

In the last section, I noted that sorting and sifting are, on an everyday understanding, properly thought of as whole-agent bodily interactions with the outside environment. Clark's (2008) account correctly acknowledges that the skills at play in a skill-based account of qualia are personal-level capacities, meaning they are possessed by the perceiver as a whole, not parts of

the perceiver. It is hard, however, to imagine what the personal-level capacities are if they are not sensorimotor skills. Clark's account is mainly pitched at the subpersonal level, and the only way to make sense of sorting and sifting being wholly inner activities is to suppose that they are carried out by some purely-inner parts of the agent, in particular the agent's brain, and become personal-level activities when the conditions are in place for these subpersonal activities to enter into consciousness. The burden of proof is on the endorser of Clark's non-sensorimotor skill-based approach to show that there are or could be any purely internal activities that can be legitimately described as sorting or sifting.

Drawing on Bennett and Hacker's (2003) view, one might argue that everyday usage is sufficient to show that 'sorting' and 'sifting' can only meaningfully describe behaviour by a whole agent. In fact, 'sorting' and 'sifting' are, in everyday practice, meaningfully applied to artifacts such as coin sorting machines and sieves, hence they fall under the limited category of intentionalistic predicates that Bennett and Hacker (2003) allow can be applied both to persons and other things. However, the Bennett and Hacker-derived point, if I were pushing it, would nonetheless apply to many of the other activities I have been subsuming under the phrase 'sorting and sifting', such as classifying, selecting, choosing and comparing. In fact, for reasons I discussed earlier, the view that everyday linguistic practice offers a definitive guide to the admissibility of applying intentionalistic or other vocabulary to a particular object is contentious, and there are apparently strong reasons to think it is false. My present line of argument does not pursue this view, but depends on the observation that sorting and sifting are as a matter of *fact*, rather than a matter of linguistic practice, commonly carried out by whole agents in interaction with the environment. All the same, any deviation from ordinary linguistic practice requires, as I argued earlier, an accompanying justification for the application of the relevant terms outside of their usual contexts.

To make a compelling case, the non-sensorimotor skill-based account must endorse the suggestion, familiar from Dennett's (1978) proposal concerning homuncular decomposition, that sorting and sifting are practical activities carried out by *subpersons* (i.e. homunculi). If this subpersonal account succeeds, subpersons could engage in their practical tasks without typically having any attendant effect on, or conceptual connection with, the whole agent's bodily interaction with the environment. However, the account must offer

compelling reasons to think these subpersons exist. Consider that when we observe the bar tender, we can tell, unproblematically, that sorting and sifting is taking place. Typically, agent-level, embodied cases like cocktail-making involve ongoing physical manipulation of the objects being sorted and sifted. Crucially, such processes must have a meaningful outcome in the sense that we can, if we like, evaluate the outcome relative to some success conditions, for instance whether or not the bar tender has assembled the correct drinks. Perception is a little different, since the perceiver need not move objects around in order to perceive. However, as with the bar tender case, it should at least be uncontroversial that we are frequently licenced to ascribe visual and/or perceptual states to agents by virtue of their outwardly observed skillful bodily engagements, leaving open the possibility that the outward bodily engagements in question constitute part of perceptual processing.

If we wanted to defend Clark's position, one thing we would have to show is that the subpersons I just mentioned can, in general, be ascribed functions, which might include sorting and sifting. One feature of cocktail-making, and of the kind of skillful bodily behaviour that allows us from the outside to ascribe perceptual states to subjects, is that it is goal-directed. To maintain in a demanding way the link between the posited subpersonal activity and the agent-level cases, the activity by the subpersons must be goal-directed too. To establish this, we would have to show not only that the subpersons happen to sort and sift, but that they have the proper functions of sorting and sifting, which is to say that they have sorting or sifting as their purpose; and the theoretical feasibility of ascribing proper functions to biological things, such as parts of the brain, is a matter of controversy. On this question, one popular line of thought has it that Darwinian natural selection offers a way to naturalise teleology, on the basis that a trait has a particular function just in case its having that function explains the trait's persistence or the organism's biological fitness. The diametrically opposed position is that natural selection - by contrast with divine creation, for example - entails that there is, properly speaking, no function or design in biology at all (Allen, 2009). Even this latter view does not, however, rule out the possibility, (i) that neural systems just so happen to perform the tasks under consideration, even if they are not proper functions, and (ii) that proper functions can be legitimately ascribed as metaphors, if not literally.

Perhaps the function issue is not of crucial importance in deciding whether or not the

ventral stream could sort and sift. For the purpose of establishing that subpersonal systems sort and sift in a way sufficiently analogous to the behaviour of whole agents, we may not actually need to establish goal-directedness in a robust way, since it is not obvious that goal-directedness is a necessary rather than contingent feature of agent-level sorting and sifting. Even if one were sceptical about proper function, then, we could posit that the ventral stream just so happens to sort and sift, with the addendum, if we wanted to endorse it, that sorting and sifting are its proper functions metaphorically speaking.

Let's take it that the aprioristic considerations allow that the idea of ascribing functions like 'sorting' and 'sifting' to neural subpersons is uncontroversial. Even where this is allowed, there persists a distinct obligation to show that there is something purely inner that can be evaluated relative to success conditions. Consider that in the first place, the homuncular subsystems require an input, i.e. something to sort or sift. These could be representations, but they could also be raw sense input signals; this is unproblematic. However, the systems must, more significantly, also have an output, some outcome by virtue of which sorting and sifting can be said to have taken place. The internalist model denies that these outputs are external to the agent as a whole, although to maintain a comparison with 'sorting' and 'sifting' in the ordinary senses of the words, the objects must be external to the homunculus.

The output of an act of sorting or sifting could, in principle, be an alteration to the same state that served as the input, as with the bar tender case, where the cocktail glasses, in one arrangement or another, feature as both the objects and the outcomes of the sorting and sifting. Since perception involves processing an ongoing stream of raw data, the homuncular subsystems that receive this input more plausibly send outputs elsewhere, i.e. to other subsystems. To licence the claim that there is sorting and sifting going on, there must be some principled way of evaluating the outputs, and this entails that the outputs must have meaning. A raw sense input signal is not intrinsically meaningful, but a representation is, by virtue being intentionally directed toward something external. This means that the non-sensorimotor, skill-based account must be accompanied by an account of representation that is both plausible and robust enough to properly account for intentionality.

6.4 - Radical Anti-Representationalism

On a 'radical' reading, the sensorimotor account is situated some way closer to the original enactivist approach that originates with Varela, Thompson and Rosch (1991), and closer still to the dynamic account of Hurley (1998), the radical embodied cognition of Chemero (2009) and the radical enactivism of Hutto and Myin (2012). The most important thing these approaches have in common is that for one reason or another they propose that perception is realised in part by extra-neural activity and takes place without internal representation. If this perspective can be convincingly motivated, either by reference to specific empirical evidence concerning perception or on broader theoretical grounds, then the sensorimotor theory can respond head-on to Clark's (2008) objection that the cognitive processing underlying perceptual awareness does depend on internal representation and not bodily movement.

Enactivism, in its best established guise, originates with Varela, Thompson and Rosch's (1991) book *The Embodied Mind*, drawing also on phenomenology (e.g., Merleau-Ponty, 1945) and earlier work by Maturana and Varela (e.g., 1980) on the biological phenomenon of autopoiesis, the process of active self-creation through which cells and larger organisms maintain their own internal structure. This tradition denies that the world, as it features in the subject's perceptual experience, results from internal representations of an objectively-specified pre-existing environment. It instead claims that the agent creates its own environment and experiential world through the same process of bodily coupling with its surroundings by which it generates and maintains its own existence. Perception, by this right, is inextricably linked to action. Even more distinctively, cognition is on this variety of enactivism inextricably linked with life (see, e.g., Thompson, 2007). The neighbouring accounts I discuss are not committed to this tradition's view that the organismic interactions that constitute perception necessarily overlap with the processes by which the organism constitutes a living thing. For this reason, they are not necessarily committed to the view that perceivers must be alive at all. However, Hurley, Chemero, Hutto and Myin, and the sensorimotor account on a 'radical' gloss, all concur that perceptual consciousness is realised by sensorimotor engagements and not representations.

Hutto and Myin (2012) adopt the 'enactive' label for their own approach, which they

call Radical Embodied or Enactive Cognition (REC). REC is motivated by the claim that 'basic' cognition, meaning any cognition not structured by natural language, 'doesn't need' and 'can't have' representational content. The authors argue that all attempts to account naturalistically for content fail because they are unable to surmount the so-called Hard Problem of Content, a deep-seated conceptual barrier to explaining how one physical state of affairs can bear truth or success conditions relating to another. The root of the problem, the authors claim, is expressed by their headline principle that Covariance Doesn't Constitute Content. This is to say that X co-varying with Y does not entail that X bears content about Y, however we flesh out the precise details of the covariance relation. For instance, a tree's rings reliably co-vary with its age and therefore carry information about its age, where information is understood as a kind of covariance. However, this does not entail that the tree trunk bears content about the tree's age, as content in this case requires truth conditions, and nothing about mere covariance entails truth conditionality. An opponent could bluntly insist, of course, that covariance *does* entail truth conditionality, to which Hutto and Myin would reply that this violates explanatory naturalism by introducing an unexplained and 'spooky' new entity.

The introduction of additional theoretical apparatus, in particular biological function, does nothing to overcome the problem so long as we continue to appeal to a covariance relation to do some of the necessary explanatory work. If Hutto and Myin's 'Covariance Doesn't Constitute Content' principle is correct, this at once undermines a large category of accounts of cognition, among which are developed attempts to account naturalistically for content (e.g., Dretske, 1981) and theories that appeal more or less in passing to the notion that covariance does constitute content (a couple of examples, which I discuss elsewhere in the thesis, include Tye, 2000; Rupert, 2011).

Hutto and Myin broaden their attack on content by targeting leading alternatives to covariance-based accounts, in particular the producer-consumer model proposed by Millikan (1984). Millikan's account suggests that content obtains in informational transactions between a 'producer' and a 'consumer', which might be different creatures or different parts of the brain. This is a variety of homuncular view, which I discussed in the last section and, earlier on, in chapter 2. Content is individuated, at least in large part, by appealing to the

purposeful manner by which the consumer responds to given pieces of information. Purpose can be accounted for naturalistically by reference to the consumer's evolved biological function. For instance, a frog's motor output system might respond to a particular input by causing the frog to dart out its tongue. If this response evolved because it allowed the frog's ancestors to eat food, then the input received by the frog's motor system carries content about food.

Hutto and Myin claim that the reference to a 'consumer' suggests that there is pre-existing content to be consumed, a violation of the principle that Covariance Doesn't Constitute Content. Where biological function does all the explanatory work, they suggest the 'consumer' would be better labelled a content 'creator'. Assuming that Millikan does mean biological function to do all the work necessary in individuating content, Hutto and Myin argue the account fails anyway. The argument, which they attribute to Fodor (1990), notes that content must involve intension as well as extension, since if there were only extension the purported content could not be false, and if it cannot be false then we cannot ascribe to it any specific set of conditions under which it is true. Natural selection, however, is only sensitive to organismic dealings with the outside environment, i.e. the extensions of the purported representations. It therefore provides no resources for explaining intension. Millikan (1993) has responded that the truth-conditions are not specified by current properties of the extension, but by properties of the extension that obtained in the historical environment where the trait evolved. However, Hutto and Myin contend that this does not solve the problem. They say: "Even if we can specify what is meant to be targeted that would give us exactly no reason to think that the targeted item is represented in a truth-conditional, referential, or otherwise semantic way - i.e. that it has intensional content" (2012, p. 80).

Their point, as I take it, is that Millikan's appeal to teleology is sufficient to show that a frog may dart out its tongue in response to a fly, or alternatively to some other object resembling a fly, using a faculty that has the proper function of responding only to flies. However, teleology is not sufficient to show that any part of the frog is evolved to bear intensional content, considering that natural selection could not be sensitive to the guise under which flies are represented, so long as the frog successfully eats them. Without intension, we could legitimately describe a frog's darting out its tongue in response to an

inedible object as an unsuccessful attempt to perform a proper function, namely to eat. But we could not make the further claim that the frog or some part of the frog represents the inedible object as food, because we have no way to specify whether it represents it as food, or as a fly, or whatever.

I take it that Hutto and Myin make a compelling case that no attempt to naturalise content has so far been given that will satisfy their metaphysical demands. One way to pursue the debate is to examine these and other attempts in further detail, and to investigate further attempts that could be made. Fruitful though this may be, a couple of striking facts undercut this. One is the sheer number of contemporary accounts featuring representation and content that appeal to nothing more than reliable covariance (e.g., Tye, 2000; Rupert, 2011). If some different or more baroque attempt to naturalise content eventually proves correct, there is no guarantee that these content-involving approaches will be able to straightforwardly substitute the new account of content for the covariance-based account while leaving all their other key claims intact. With this in mind, Hutto and Myin will have made substantial progress toward making radical enactivist approach to cognition appealing if they can successfully unseat covariance-based accounts of content.

Offsetting this argumentative advantage is the possibility that Hutto and Myin are simply making unduly grand metaphysical demands of their opponents. Although Hutto and Myin probe the specifics of individual attempts to naturalise content before arguing that they fail, their position is not altogether different from Bennett and Hacker's (2003). Hutto and Myin, for their part, assume without much argument that the meaning of words including 'represent', 'content' and 'truth condition' are such as to place a demanding metaphysical burden on anything that can be correctly said to fall under those terms, and go on to examine specific proposals before finding that they are unable to meet this metaphysical demand. Bennett and Hacker, by contrast, assume without much argument that there can be no underlying feature shared by persons and parts of persons that might actually licence the ascription to both of predicates like 'representing'. Where Bennett and Hacker make a linguistic demand, they appeal implicitly to a metaphysical assumption. Hutto and Myin make a metaphysical demand, but appeal implicitly to similar linguistic assumptions. One theorist who disputes this approach is Chemero (2009), whose 'Radical Embodied Cognitive

Science' defends a similar positive programme to Hutto and Myin's, but who sharply disagrees with their negative claims.

Chemero argues that in line with the linguistic conventions employed by scientists, the words 'representation' and 'content' make only modest metaphysical demands. By this latter standard, certain kinds of covariance relation do, in fact, constitute content. Hence, we could well say that a tree's rings represent or bear content about its age. Defending content-free cognition, Chemero more modestly insists that cognitive scientists should only appeal to representation or representational content if doing so provides the best explanation for a given phenomenon. This, he takes it, is an 'epistemic' rather than a 'metaphysical' demand. Hutto and Myin say something superficially similar when they concede that the appeal to representational content would be legitimate if the explanatory success of perceptual science depended on it. However, unlike Chemero, Hutto and Myin suggest that even where the appeal to content is motivated by epistemic rather than metaphysical considerations, it nonetheless entails a robust metaphysical commitment of the sort Chemero rejects.

In one way or another, each side justifies its stance by appeal to naturalism. For the present purpose, we may take it that the naturalism in question is nothing more robust than the minimal naturalism advocated by Wheeler (2005) and endorsed by Hutto and Myin. This amounts to the claim that science and philosophy must be consistent, and that where their respective claims conflict, philosophy must give way. Chemero argues in effect that where scientific accounts find theoretical utility in positing 'representations', and apply the word in a manner consistent with their own linguistic convention, philosophical considerations cannot unseat them, since philosophy must give way to science. In this sense, Chemero is assuming the mantle of explanatory naturalism.

Hutto and Myin deploy minimal naturalism to the opposite effect. They take it that notwithstanding the apparently productive deployment of the terms 'representation' and 'content', especially in empirical work, no present theory shows how the physical states in question genuinely represent or bear content in light of the metaphysical demands Hutto and Myin think apply. Their point, to extrapolate a little, is this. Explanatory utility is insufficient to licence the claim that some entity exists if there is otherwise no empirical or theoretical evidence that the entity does, or even could, exist. If there is no reason to think that an entity

exists, it is unscientific to claim it does exist. On this basis, the authors appear to have naturalism on their side.

In Chemero's favour, it is evident that scientists, especially neuroscientists, frequently use 'representation' and even 'content' to denote nothing more demanding than covariance. Moreover, it is true that people are entitled to use words in accordance with the linguistic conventions in operation within their own discipline. In Hutto and Myin's favour, the problem is that 'content' is also frequently used to denote something metaphysically more robust than mere covariance, even on accounts that appeal to nothing more than a covariance relation to account for it (Tye, 2000, is a good example, since his account of phenomenal qualities depends strongly on the idea that they reduce to representational content and not mere co-variance relations). Where content is used to denote a covariance relation, it is, given this equivocation between speakers, not always apparent whether the speaker means to make an extravagant metaphysical claim about the covariance relation in question, or means to use 'content' in a deflationary way to denote something metaphysically austere.

To resolve this impasse, both sides ought to endorse McDowell's (1994b) position that where content is ascribed subpersonally, i.e. to parts of persons rather than persons, it should be understood by default, and in all cases, as no more than a metaphorical ascription. 'As if' content, as he calls it, is not really content, but may be legitimately ascribed just in case it proves useful or indeed indispensable as a theoretical posit to explain information processing transactions between an agent's parts, a possibility he happily endorses. By adopting this as a permanent disclaimer, it is possible to accommodate Chemero's position that representation and content are legitimate posits when deployed in a metaphysically austere way. At the same time, this ought to resolve any suggestion that 'content', deployed in a superficially harmless way, actually disguises a more robust metaphysical commitment.

Although metaphorical content is admissible in principle, it is subject to certain provisos. One is that if we aim to account for phenomenal qualities naturalistically in a way that defeats anti-physicalist intuitions, we cannot do so by appealing to solely to metaphorical content. Although a physical state that merits the metaphorical ascription of content may be real, and the ascription of content genuinely useful for some purposes, the content itself, if metaphorical, does not really exist; it is just a way of talking. So if you think that a 'red' quale is

a non-physical state, you have no reason to change your mind on discovering that it is correlated with metaphorical content. However, if perception is constituted by bodily interactions with objects in the environment, we can appeal to the qualities of the interaction to account for phenomenal qualities, and use metaphorical content, if appropriate, to explain how those interactions take place. This brings us to the other proviso, however, which is that content and representations may not be the most appropriate way to account for the physical interactions involved in perception, as we will see in the next section.

6.5 - Representation Restricted

According to Marr's (1982) theory, and accounts in the same tradition, conscious perception is realised by the brain's construction and deployment of representations. This processing is properly understood as an enabling rather than constitutive feature of perception. A constitutive account of perception says what perception is. An enabling account, in this context, could be thought of as a variety of causal account, where we take 'A causes B' to imply that A happens at the same time as B but at a lower explanatory level. Endorsers of the personal/subpersonal distinction observe that the representations in Marr's account must be enabling rather than constitutive features of perception because they are subpersonal, and perception is a personal-level faculty (McDowell, 1994b; Pessoa et al., 1998). Noë (2004) denies that the personal/subpersonal distinction is important, which means he cannot so straightforwardly establish that subpersonal representations are not constitutive features of perception. However, like Pessoa et al, he makes use of Marr's own distinction between the task and representational levels of analysis. The former specifies which tasks perception aims to perform, while the latter describes how the brain processes information, in the form of representations, to enable these tasks to be performed. As Noë appreciates, the distinction between the task and representational levels should be taken to track the constitutive/enabling distinction, meaning representations are mere enabling features even if the task level is not necessarily the personal level.

McDowell suggests that by appreciating the constitutive/enabling distinction, a rapprochement can be achieved between Marr's representationalism and ecological perception,

which is non-representationalist. His claim is that a Marr-style representationalist account can usefully be given of the subpersonal processes that enable perception even while we use the ecological approach to characterise perception's constitutive features at the personal level. As Pessoa et al. and Noë (2004) bring to our attention, however, Marr has a faulty view of representation that prevents this reconciliation from being achieved. The problem arises because Marr depends too heavily on his representational level framework to inform his constitutive or task level account of vision. Indeed, he appears to straightforwardly conflate the enabling with the constitutive features when he claims that: "vision is the *process* of discovering from [retinal] images what is present in the world, and where it is" (Marr, 1982, p. 2). The bluntest objection we could make to this is that it commits a straightforward category mistake. As McDowell has argued, vision is personal level process in which we come into contact with the world rather than the retinal image. Pessoa et al. are sympathetic to this point, but their argument does not focus on the bare coherence of claiming that perception is constituted by a subpersonal process. Their criticism is mainly directed at the danger that a subpersonal constitutive account invites implausible commitments including, as we saw in chapter 2, the undue commitment to a subperson, a homunculus, for whose benefit processes like filling-in need to take place.

Noë (2004), having denied that the personal/subpersonal distinction is important, has no call to insist that a constitutive account of vision must make reference to the whole person (or more broadly, agent). His relatively nuanced criticism, reflecting a point originally made by Nakayama (1994), is that Marr fails to appreciate the difficulty of giving an appropriate task level (i.e. constitutive) account of vision. As I take it, the problem is that Marr attempts to give an independently plausible account of perception's representational (i.e. enabling) features, and lets his view of the task level be unduly constrained by it, a bias implicit in his claim that perception *is* a process in which we use the retinal image to find out about the world. This is a mistaken order of priority, because our understanding of the processes that enable perception ought in many cases to be substantially constrained by independently motivated constitutive claims. Instead of only allowing one-way traffic from the enabling to the constitutive, or for that matter in the other direction, there should be a two-way interplay between our understanding of perception and our understanding of the processes that enable

it (see, McDowell, 1994b; Wheeler, 2013).

The objection is not that Marr's account, and Marr-style accounts, are framed as if no possible evidence, in particular empirical evidence, could disprove them. The problem is that their apparent plausibility as enabling accounts means that should they unduly inform or be conflated with constitutive accounts, they close off the possibility of more carefully examining perception's constitutive features. Re-emphasising neural representation's rightful place only as an enabling feature has a useful dual effect. It opens up conceptual space for a better constitutive account of perception. This, in turn, opens up the possibility that we must re-evaluate our claims about the processes that enable perceptual consciousness. This could have various effects, for instance constraining in previously unrecognised ways the kinds of content or representational vehicle that plausibly enable conscious perception, and, in principle, making it plausible that representation should be partially or even wholly dispensed with in favour of appeals to dynamic bodily interaction with the outside environment.

A proper constitutive account of perception ought to draw on a number of distinct methods and considerations. This includes empirical work. For example, the work on change blindness and inattention blindness referred to by O&N is enough to strongly indicate that the brain does not represent, all at once, a large and richly detailed visual field. A constitutive account should also pay careful attention to the phenomenology of visual experience, an aspect of sensorimotor theory most prominent in Noë's (2004; 2012) accounts. Noë claims that visual experience does not feel picture-like. This is, as I think, a straightforward and compelling aspect of Noë's in some ways enigmatic description of visual experience, exemplified by his claim that: "A perceptual experience doesn't analyze or break down into the experience of atomic elements, or simple features [...] the moment you stop and try to make a specific feature the sole object of your consideration – this shade of red, for example – it slips away from you in the sense that it exceeds what you can take in, in completeness, in an instant" (2004, p. 135).¹⁷ The fact that visual phenomenology is not pictorial constrains the shape that an enabling account ought to take. It offers one reason, for instance, to doubt that the character of perceptual awareness should be explained by pictorial representations in the

¹⁷ I offer a more detailed take on this point in chapter 7.

brain.

Noë suggests that attending to one's own visual phenomenology is enough to demonstrate that perceptual consciousness, instead of being pictorial, is constituted by the exercise of Sensorimotor Understanding. This claim has some appeal, but is hard to make go through purely on the basis of attention to one's own experience. A further argument is made by Noë (2010), who claims, as we saw earlier (section 4.3), that a perceiver would not be able to integrate what they consciously see with what they do if conscious vision did not involve Sensorimotor Understanding, or a similar egocentric understanding of their bodily alignment with objects in the outside environment. This can be understood as a logical, constitutive claim about vision, and places downward pressure on our understanding of the processing carried out in particular by the ventral stream to realise visual awareness. Notice that the point comes through even more sharply, however, when we think of it as an engineering question, namely the question of how the ventral stream achieves its function of helping the dorsal stream produce appropriate skillful bodily behaviour. The idea that the ventral stream must trade in representations of SMCs is a low-level claim about the processing that enables perceptual consciousness which we find, it turn, places upward pressure on a high level account of conscious perception's constitutive features. This is one case in which the sensorimotor theory makes productive use of an interplay between explanatory levels.

On the view of Sensorimotor Understanding I endorse, in which it is a feature of the perceiver's behavioural capacities, the claim that Sensorimotor Understanding is a constitutive feature of conscious perception could be supported by the need for ventral stream processing to integrate with dorsal stream processing too, especially if we endorsed the view that the ventral stream integrates with dorsal stream activity by representing possibilities for coarse-grained action (Ward, Roberts and Clark, 2011). Sensorimotor Understanding, on the behavioural capacity view, also makes a good fit with non-representationalist, dynamical systems accounts of perception, considering that they already appeal to sensorimotor capacities (Hurley, 1998; Noë, 2004; Chemero, 2009; Silberstein and Chemero, 2012; Burhmann, Di Paolo and Barandiaran, 2013).

A constitutive account of vision should also take account, I argue, of philosophically motivated worries about intentionality and phenomenal consciousness. Noë (2004) attempts

to do justice to McDowell's (1994a) claim that the content of perceptual experience is necessarily conceptual by suggesting that the bodily capacities we draw on in conscious perception are themselves conceptual. Noë (2012) retracts this claim and instead states that sensorimotor skills involve a form of nonconceptual content, while occurring on a continuum with conceptual skills rather than being radically distinct. I will not comment on the debate between conceptualists like McDowell and his opponents (e.g., Dreyfus, 2013; Evans, 1982), who claim that action and perception can be intentionally-directed without being fully conceptual. However, it is important to do justice to the basic insight that perceptual consciousness must be intentionally directed, and that this requires the perceiver to bring to bear some kind of skillful, active engagement with the world. It will not do to assume that content borne by neural representations can play this role, since it has not been convincingly established that neural representations bear genuine rather than metaphorical content. For this reason, appealing to neural states alone will not do at all. Accounts featuring content-bearing neural representations plausibly provide, in some cases, a useful framework for explaining the role the brain plays in producing behaviour, but it has not been established that they can, except by reference to behaviour, help account for genuine intentionality.

We might be tempted to deny that a naturalistic account of cognition needs to answer Hutto and Myin's metaphysical worries about the ascription of 'content', or that it is the job of a scientifically-oriented account of vision to explain how the capacities appealed to by non-naturalistic accounts of perception can be accounted for in physicalistic vocabulary. Consider, however, the purpose of the skill-based approach to phenomenal qualities, endorsed by Clark (2008) and on Clark's reading, with which I agree, by the sensorimotor account. The skill-based approach is meant to avoid the intuition that phenomenal consciousness is non-physical by identifying phenomenal consciousness with the exercise of skills. Its job is to avoid the dilemma I mentioned earlier (adapted from Ryle, 1979, by Silberstein and Chemero, 2012) in which phenomenal consciousness is either 'nothing but' brain activity or 'something else as well'. As I have argued, if the skills appealed to by the skill-based account are not sensorimotor skills, they depend on internal representation. If that representation is metaphorical, or a useful fiction (Sprevak, 2013), or anything other than robustly real, it cannot help answer metaphysically motivated, anti-physicalist intuitions about consciousness

better than appeals to nothing but brain state. This, in my view, is the biggest problem with attempts to defuse non-physicalist intuitions about qualia with appeal to representation.

One of the best reasons for taking Sensorimotor Understanding to be constituted by a capacity for skillful behaviour - rather than merely by knowledge of the sensory consequences of possible movement - is that it offers a promising way to account naturalistically for the intentional directedness, and by virtue of this the phenomenal character, of perceptual consciousness. Giving a proper account of the criteria a repertoire of behavioural capacities must meet to count as constituting mindedness, and sensorimotor understanding in particular, is not something I pursue here, and is something that would eventually need to be provided for this point to be established. However, it is hard to view the outward behaviour of a human perceiver and resist the intuition that they are intentionally-directed toward the features of the environment they interact with. This provides one motivation for regarding behavioural capacities as constitutive features of perceptual consciousness.

If we could be certain that neural representations cannot be genuinely contentful, then appeal to sensorimotor skills would appear to be the only hope for finding a naturalistic account of intentionality and so phenomenal character, and the case for regarding them as constitutive features of perception would be secure. As we have seen, Hutto and Myin offer strong reasons to doubt that subpersonal content is naturalisable by appeal to covariance or teleology. A possibility Hutto and Myin endorse, which I also endorsed in section 2.1.2, is that content could be adequately naturalised by appeal primarily to empirical evidence. In principle, it might be possible to routinely provide full physicalistic descriptions of the neural processes that underlie a conscious experience, and use a subpersonal theory that makes essential explanatory appeal to representation to provide bridging principles which make the relation between the neural processes and the conscious state intelligible. Without a conceptual solution to the 'hard problem' of content, it would remain hard to robustly justify the claim that neural states bear genuine rather than make-believe content, or in the language of McDowell (1994a), to say that the brain is responding in a person-like way to reasons and not just causes. However, the empirical success of the enterprise just described would provide strong grounds for abandoning the requirement for a naturalistic theory of perceptual consciousness to meet these strong metaphysical demands.

That such an account could be provided is an empirical bet. Although it is possible to provide representation or content-involving accounts of neural activity that provide an explanatorily useful framework for understanding brain activity itself, there is no guarantee that this subpersonal content can correspond reliably to the content or phenomenal character of perceptual consciousness (Noë and Thompson, 2004). The problem is not the epistemically-founded worry that it is difficult to intelligibly describe and explain the neural processes that yield perceptual consciousness unless we make reference to the contribution they make to sensorimotor coupling. The problem is that there may be no intelligible link between purely inner processes and perceptual consciousness. This point is evident in Hurley's (1998) defence of vehicle externalism, although we could accept the modest conclusion that an account of perceptual consciousness ought to make explanatory appeal to sensorimotor coupling without necessarily insisting on ECM, the claim that the substrate of perceptual consciousness actually includes the body and outside environment.

The question of whether perceptual consciousness depends on internal representation or a non-representational dynamical system depends on how well the respective approaches make sense of empirical work. However, there are aprioristically motivated reasons for not being surprised if it turns out to be impossible to explain perception by attributing content to brains or parts of brains, namely that intentionalistic vocabulary like 'believing' or 'representing' are words we apply in the first place to whole persons and not their brains. The fact that strategies like homuncular decomposition make it conceptually viable to apply these predicates to brains does not mean that we should expect these strategies to finally allow us to explain conscious perception by appeal only to contentful states of the brain. A better bet is to regard bodily capacities as perception's constitutive features and use these as the starting point to account for the physical processes that enable perceptual awareness.

7 - Object Experience as Temporal Experience

This thesis has so far focused largely on defending and clarifying the basic tenets of the sensorimotor theory. With these tenets shored up, there is as-yet untapped potential for the sensorimotor theory to explain in much greater detail the distinctive contribution it can make to accounting for specific perceptual phenomena. For instance, there is much work to be done on identifying the sensorimotor contingencies involved in colour perception, and the theory has not yet given much account of sense modalities other than vision. This present chapter continues the work clarifying and defending the sensorimotor theory, but also sketches out a novel contribution the sensorimotor theory can make to a long-discussed puzzle concerning the experience of temporally-extended events. In addition to explaining how the sensorimotor theory might make sense of event-experience, I will argue that object-experience may itself be a variety of event-experience, and that this would help make sense of Noë's appealing but oblique account of visual phenomenology.

7.1 - The Puzzle of Temporal Experience

The puzzle of temporal experience (Kelly, 2005) concerns perceptual awareness of temporally-extended phenomena like change, succession and constancy. To illustrate, Kelly offers the example of an opera singer delivering a sustained note. When you hear it, your experience incorporates not just the note's immediate presence, but also, somehow, its extended temporal duration:

There you are at the opera house. The soprano has just hit her high note – a glass shattering high C that fills the hall – and she holds it. She holds the note for such a long time that after a while a funny thing happens: you no longer seem only to hear it, the note as it is currently sounding [...] in addition, you also seem to hear something more [...] the note now sounds like it has been going on for a very long time [...] What you hear no longer seems to be limited to the pitch, timbre, loudness and other strictly audible qualities of the note. You seem in addition to experience, even to hear, something about its temporal extent. (Kelly, 2005, p. 208)

Vision, too, sometimes comprises the experience of temporally-extended events. Grush (2007)

offers an example of motion experience. Suppose you are looking at a clock. You could stare at the hour hand all day, and although you may infer that it is moving, you never visually experience its motion. When you look at the second hand, however, you seem to perceive its motion directly as it moves around the clock face. Temporal experience is a puzzle, as Kelly (2005) puts it, rather than a mere problem, because it involves an apparent paradox: What you perceive now, at this moment, ought only to be what is present now. Motion, change and duration are not momentary, but take place over extended periods. Solutions to the puzzle will differ in their diagnosis of what is happening when you hear the opera singer, or look at the second hand, but the examples clearly illustrate that there are data to be explained.

Dainton (2010) places proposed solutions into three distinct categories. The 'cinematic' approach is the view that perceptual experience is divided into static snapshots, like a cinema reel. On this approach, what you perceive at a given instant is just what is present at that instant, meaning you cannot right now literally experience a temporally-extended event. Temporal experience, here, has to be explained some other way. Crick and Koch (henceforth C&K) (2003), for example, suggest that you do not experience motion, but "a series of static snapshots, with motion 'painted' on them" (p. 122); they suggest, by analogy, that the experience of motion has something in common with the motion suggested by a drawing of a person in mid-stride (p. 122, figure 1). The other two types of explanation, by contrast, endorse the existence of what James (1890/1981) called the 'specious present', an experiential 'now' that really does have duration.

The 'retention-protention' approach (Dainton, 2010), of which Husserl (1991) and, under Husserl's influence, Varela (1999) were notable exponents, says that the experience of what is genuinely present at a given instant is accompanied by 'retentions' from the immediate past and/or 'protentions' into the future; these can be understood, respectively, as special perceptual sorts of memory and anticipation. Grush (2007) advocates a retention-protention account he calls the Trajectory Estimation Model. In support of the approach, he cites empirical work on a number of illusions, for example the 'cutaneous rabbit' illusion (Geldard and Sherrick, 1972, cited by Grush, 2007). The study involved subjects being given a sequence of taps on the arm, five at each position, in three successive positions along the arm. The subjects reported that instead of feeling the taps in just three different positions, they

experienced them as travelling in an even trajectory along the arm, ultimately extending beyond the location of the final tap. Because this trajectory would not have been experienced had the sequence been confined to the first five taps, Grush's suggestion is that an experience at a given instant can feature as part of its content prior occurrences (see also the colour-phi phenomenon, discussed, e.g., Dennett, 1991). Further, because the felt trajectory continued onward, Grush concludes that protentions – anticipations of future events – also figure in experiential content.

The 'extensional' approach (or 'extensionalism') – endorsed, for instance, by Dainton (2000) – says that the content of the perception (the state of affairs represented by the experience) temporally tracks the vehicle (the physical state that realises the experience). Dainton (2010) observes that on this view, since “our episodes of experiencing are themselves temporally extended, [they] are thus able to incorporate change and persistence in a quite straightforward way”.

Clark (2006) uses temporal experience – in particular, Kelly's opera singer example – to expose what he regards as a general problem with the sensorimotor theory. His argument asserts that temporal experience poses a problem, in particular, for Noë's notion of presence-as-access, the view that you fail to experience the objects of perception as straightforwardly 'present', but instead experience them as 'accessible', thanks to your possession of Sensorimotor Understanding. In targeting presence-as-access, Clark implicitly targets the claims that perceptual experience depends on Sensorimotor Understanding and the claim that phenomenal character can be characterised by reference to SMCs. Clark says we cannot explain the experience of the opera singer's note “by appeal to any sense of the potential availability of the missing parts of the temporally-extended sound stream, nor can we know (indeed, it is barely intelligible to ask) how those missing parts of the soundstream would vary or come into focus as we move our head or body” (Clark, 2006, p. 23). Since the past and future are not, in other words, mediated now by laws of sensorimotor contingency, it cannot be Sensorimotor Understanding that explains the experience of the note's duration.

Noë (2006) responds by agreeing that Sensorimotor Understanding only explains object perception, arguing that this does not compromise the sensorimotor theory in general. He suggests that since event perception is, in any case, a quite different species to object

perception, the theory is not committed to giving a sensorimotor account of the opera singer case. He justifies this by reference to the distinct qualitative character of the temporal experience compared to the experience of an object: “it rides roughshod over the phenomenology [...] to say that the past sounds [like objects] are now present or that they are now accessible” (p. 28); instead, he claims, you hear the note “as having a certain trajectory or arc, as unfolding in accordance with a definite law or pattern” (p. 29).

Noë’s positive account of the opera singer case posits that we have a (non-sensorimotor) grasp of where the note is coming from, analogous to linguistic understanding:

When you hear the singer’s sustained note, you do not experience the acoustical properties of the sound, any more than you experience the acoustical properties of the words you hear when you understand speech. In the linguistic case, you hear meanings themselves, you hear what is said. In the case of the singer, what you actually hear is the singer herself, her voice, her vocal action – what she is doing. It is the fact that the singer is doing something, performing an action, that fixes the relevant temporal horizon and intentional arc (Noë, 2006, p. 29)

This offers a solution to the temporal experience puzzle akin, in important respect, to C&K’s cinematic account. Where Noë claims that different kinds of knowledge are responsible for object experience and temporal experience, C&K, similarly, ascribe these to separate “mechanisms” (p. 122). Moreover, both C&K and Noë agree that you do not experience temporally-extended phenomena directly: As with Noë’s invocation of an experienced trajectory, C&K suggest motion is “painted” on” (p. 122) to static snapshots, resulting in your experiencing moving objects as being in motion, while failing to directly experience that motion.

However, Noë maintains that even temporal experience involves a coupling with the environment rather than a representing:

[P]erception is an activity of sensorimotor coupling with the environment [...] experiences are not acts [...]; they are not representations; they are activities, events themselves; they are temporally extended patterns of skillful engagement. When you perceive an event unfolding, it is not as if you occupy a dimensionless point of observation. You live through an event by coupling with it. (Noë, 2006, p. 31)

This suggests a quite different line of response than the one endorsed by the analogy with linguistic understanding. Extrapolating from Noë, the thought seems to be that the physical substrate of the experience is a smoothly continuous activity rather than one which breaks down into temporally discrete chunks: hence to explain the physical substrate of perception, one must look at dynamically unfolding interactions, rather than ‘object’-like structures in the brain. If we appended to this the extensionalist view that the content of an experience temporally tracks the activity of experiencing, it would mean that the content of experience is always, itself, temporally-extended. Notably, this does not need to entail that past and future portions of the note should be experienced as now present. Rather, it suggests that to experience the world is – as Noë claims – to experience being in the midst of some trajectory.

7.2 - Coupling and Knowing in Temporal Experience

The argument over temporal experience speaks to a broader tension within the sensorimotor account. As many commentators have noted, the approach sometimes appears to stipulate that the temporally-extended bodily exercise of sensorimotor skill is required; at other times, it apparently suffices that the perceiver possesses sensorimotor skill. Aizawa (2010) describes these as ‘strong’ and ‘weak’ versions of the theory. Noë’s discussion of temporal experience betrays a similar tension: while its main claim is that the experience of duration is explained by your non-bodily understanding of what you are hearing, it indicates a quite different line of response when it argues that hearing the note involves, in any case, a temporally-extended coupling with the environment. The tension is particularly stark in the case of event perception, because it is hard to see why temporal experience should be explained both by understanding and sensorimotor coupling. In the original story about object perception, sensorimotor knowledge is sometimes glossed as knowing how to act, for example when O&N suggest sensorimotor knowledge is comprised of “action recipes” (p. 945) or when Noë says: “[t]o experience [an object] as on the left is to experience it as necessitating [...] various possibilities of sense-affecting movements” (2004, pp. 87–88). Here, we can readily grasp that perception might involve, vitally, a temporally-extended process of bodily coupling, itself

featuring Sensorimotor Understanding. By contrast, your knowledge of where the opera singer's note is coming from may be implicit, but there is no obvious way in which it is practical, or geared toward action, as sensorimotor mastery might be. If it is not practical, but, as Noë suggests, more like linguistic comprehension, then it is hard to see why possessing or exercising this knowledge should entail, in any interesting sense, a coupling with the environment rather than just a representing of the environment. The sensorimotor theorist, in the temporal case, no longer has an obvious response to the theorist who maintains that neural states or structures alone are the interesting, indeed constitutive, features underlying conscious perceptual experience.

This threatens to undermine the sensorimotor theory, as it means the 'weak' variant, on the current account of temporal experience, can no longer sustain any of the theory's main tenets. Vehicle externalism or ECM is ruled out, since bodily movement is not required for perceptual experience. Sensorimotor Understanding, SMCs and presence-as-access only give an incomplete account of perceptual experience because they account only for object experience, not temporal experience. The sensorimotor theorist could respond that although event perception is non-sensorimotor, it depends for its existence on object perception, which is sensorimotor. Conceding this much, however, gives the opponent room to deny that perception is intrinsically sensorimotor at all: Clark (2008), as we have seen, claims that perception is a matter of sensorimotor summarising – the extraction of information about sensorimotor contingencies, along with other information, for the construction of representations that are not themselves finely sensitive to the sensory effects of possible movements.

Moreover, the distinction made by Noë between events and objects seems tenuous. Noë (2006) says "objects [unlike events] are timeless in that they exist whole and complete at a moment in time" (p. 28). The implication is that it is therefore reasonable to suppose that *perception* of objects and events are different matters. On this view, perception of a moving object would seemingly be a matter of perceiving the object (sensorimotor) and perceiving its trajectory, an event (non-sensorimotor). However, if you perceive an unmoving, unchanging object, and perceive it to be constant, you expect that it existed moments before, and that it will continue to exist moments into the future. As a result, it is not enough to tack on some

kind of non-sensorimotor perception just for some instances. The non-sensorimotor element would apply to all instances of perception.

The opera singer example is a little misleading in this respect, because the phenomenal experience of the note's temporal duration seems like an optional extra, a side issue to the perception of immediately audible qualities like pitch, loudness, and so forth. Temporal aspects to perception are not usually like this. When you cross a busy road, or catch a ball, the perception of objects as still or in motion is no side issue, but central to your engagement with them. For this reason I think it is better to hold that object and event perception are not different types. The fuzziness of the distinction between object and event perception provides another reason to prefer a unified picture for objects *and* events, and conceding that events are not perceived in sensorimotor fashion provides a reason to believe that objects are not perceived in such a fashion either. This is one reason why I will argue presently that a unified sensorimotor model can be offered for both.

The best solution, I suggest, is to drop Noë's analogy with linguistic understanding and stick, instead, to an extensionalist story, which says that event experience is explained by a particular kind of temporally-extended coupling with the environment. This requires adopting the 'strong' sensorimotor theory, in the sense defined above, and accords with my suggestion in chapter 4 that Sensorimotor Understanding is a criterion that temporally-extended engagements must meet, rather than an entity that enables the appropriate engagements to take place. The mere possession of sensorimotor skill would not be sufficient to explain the experience of duration, since duration is not a matter of sensorimotor contingency. However, the bodily exercise of Sensorimotor Understanding takes time anyway: so, if we assume the content of experience temporally tracks the vehicle, the experience of duration comes for free. This suggestion is not merely a get out clause, but fits the phenomenology – the experience of trajectory – aptly described by Noë. It explains why your experience, now, of the opera singer's note sounds like it is part of something temporally-extended.

7.3 - Object Experience and Skill-Based Access

I have suggested that the best response the sensorimotor theorist can make to Clark's objection from temporal experience is to argue that the experience of temporal duration supervenes on a temporally-extended physical process, comprised of skill-driven bodily coupling with the environment. This means that the exercise of Sensorimotor Understanding must be temporally-extended in the case of event experience. The sensorimotor theory could, conceivably, endorse this, but claim that object experience does not depend on temporally-extended engagements, merely on the activation of representations of SMCs. In this case, the debates over temporal experience and what Aizawa calls 'weak' and 'strong' variants of the sensorimotor theory come apart. My proposed version of extensionalism, however, accounts for object perception as well as event perception. This results from what I contend is the most productive way to understand presence-as-access, and it is this thesis I turn to now.

As we saw earlier, Noë (2004) outlines his notion of skill-based access when he likens perception to accessing a newspaper via the World Wide Web. He points out that when you view the online version of the New York Times, your computer does not download the day's edition all at once, but downloads, on request, one article at a time. This is sensible, as it limits the burden placed on your computer and internet connection, and means that should an article be updated, you get the latest version. Importantly, accessing the paper in this piecemeal fashion is, for all intents and purposes, just like having the whole issue there at once, since every article is accessible as needed. The day's edition is available, as Noë puts it, *virtually*.

Seeing, he suggests, works in a similar way. You do not experience, all at once, a richly detailed visual field. Instead, you access, as required, detail available from the outside environment. The presence of rich detail is also, in this sense, virtual (Noë, 2004, pp. 49–51). However, in a crucial disanalogy with the computer case, Noë claims that perceptual presence is "virtual *all the way in*" (2004, p. 134). The suggestion is apparently that no complete property is ever present to experience in a local or offline manner: a claim Noë argues is justified by attention to one's own experience:

A perceptual experience doesn't analyze or break down into the experience of atomic elements, or

simple features [...] the moment you stop and try to make a specific feature the sole object of your consideration – this shade of red, for example – it slips away from you in the sense that it exceeds what you can take in, in completeness, in an instant. (2004, p. 135)

This is a puzzling sort of claim. There is (I think – and the reader might accept, at least for the sake of argument) something in Noë’s phenomenological claim that you cannot, in a momentary act, get a visual grasp on any complete property, such as a shade of red. However, there is certainly something it is like to see a red thing. Noë explains this by claiming that properties are “present not as represented, but as accessible [...] [t]hanks to my possession of sensorimotor skills” (2004, p. 215). As he puts it in a later piece, the environment:

[...] shows up as present, but out of view, in so far as I understand that I am now related to it by familiar patterns of motor-sensory dependence. It is my basic understanding of the way my movements produce sensory change given my situation that makes it the case, now, even before I have moved an inch, [emphasis added] that elements outside focus and attention can be perceptually present. (Noë, 2012, p. 19)

This view is a ‘weak’ position, since it stipulates that no movement or temporal extension is required. As a result, the perceptual experience described must be the result of an internal state, perhaps an internal representation specifying what movements the perceiver should make to bring objects into view. It is not obvious, however, why this sort of Sensorimotor Understanding should not allow a perceiver at some instant to take in an atomic visual feature, given visual features are meant to be specifiable by reference to movement-related contingencies. We could make sense of this by supposing that the sensorimotor mastery made use of by the perceiver at a given moment fails to correspond, precisely, to any atomic feature, such as a shade of red. If this were true, however, it would not be clear how we can explain the fact that perceivers experience shades of red at all. As I result, I contend that the best way to make sense of the computer metaphor, and the associated phenomenology, is to hold that object experience, like event experience, depends on a temporally-extended physical process. Indeed, the extensionalist proposal is useful as a way to characterise both event and object perception. I elaborate on this in the next section.

7.4 - Extensionalism About Event and Object Experience

It is beyond my present scope to consider the advantages and disadvantages of various approaches to temporal experience independent of their relation to sensorimotor theory. I argue, however, that extensionalism is likely to provide the right conceptual foundation for a sensorimotor account of temporal experience. Adopting this approach is useful for the sensorimotor theory, because it allows the theory to respond convincingly to the specific challenge set by Clark (2006), and thereby save the approach from the danger that it fails altogether because it is unable to account for temporal experience. More significantly, there is good reason to endorse the extensionalist claim that perceptual experience temporally tracks a temporally-extended physical process anyway, since it has the added utility of offering a new and better way of explicating some of the compelling claims featured in Noë's account of presence-as-access. On my proposal, event and object experience are not only explained the same way, but are aspects of the same phenomenon.

An extensionalist sensorimotor account, in the first place, explains temporal experience by claiming that perceptual awareness of a temporally-extended event supervenes on a temporally-extended process of interaction between the perceiver and the environment, in such a way that the content temporally tracks the vehicle. This is, in principle, compatible with the view that object experience is explained by the mere activation of a representation bearing content about SMCs. However, we could plausibly go further and suppose that the characterisation of object experience found in Noë's claims about presence-as-access can best be explained by accounting for object experience in the same way as event experience.

An initial move is to reject any suggestion that knowing the movements you could make to bring an object into view can amount, by itself, to entertaining perceptual content about that object. Instead, knowing this might enable, more modestly, a nonspecific feeling that there is detail available which can be accessed from the environment as needed (as suggested by Clark, 2008, p. 194). To account for perceptual presence – of an opera singer's note, or a shade of red – I suggest we instead point to the perceiver's skillful coupling with an environment in which the relevant laws of sensorimotor contingency apply. Suppose, for the sake of argument, that this is a correct account of object experience. Combined with the

extensionalist claim that the content of perceptual experience temporally tracks the physical substrate, it would follow that the content of object experience is itself temporally-extended. This is counterintuitive, but compelling if you consider the presence-as-access in a particular light. Recall two of the claims featured in this aspect of the sensorimotor account.

There is a phenomenological claim, concerning a perceiver's inability to experience, in a momentary act, any atomic visual features. There is also a claim about the material substrate of experience, captured by Noë's discussion of virtuality, which suggests that perception is entirely beholden to the agent's online interaction with the environment. Neither of these claims is indisputable, but they are coherent accounts of how perception may work. If you think there is something right about the theses, extensionalism about object experience offers a way to make sense of it. On the view I recommend, one cannot, in a momentary act, take in a shade of red for the very same reason that one does not experience, at this instant, the past and future portions of the opera singer's high C. Perceptual experience, both of objects and events, supervenes on something that may, essentially, be understandable only by reference to a temporally-extended, dynamical process.

Noë's (2004) claim that experience is "virtual *all the way in*" (p. 134) could suggest that the spatial size of the visual field is infinitesimally small or even non-existent. This makes the sensorimotor account sound like it denies perceptual experience altogether. A better gloss on the claim says that experience is virtual all the way in not, primarily, because the information processed at one time by the visual system is minimal – although it may be minimal – but because this information can only be processed in the course of temporally-extended coupling with the outside environment. To revisit the World Wide Web analogy: if your access to the online newspaper were virtual all the way in, a faulty understanding of virtuality would suggest that your computer 'knows' how to access the internet, but has no screen with which it can locally display a web page. According to my gloss, however, the computer has a screen, but loses its ability to display any web page – even one you already have open – the instant it loses its connection to the Internet. It is a commonplace that perception, ordinarily, requires that there is an environment present to perceive. However, the point of virtuality, as currently understood, is to show that ongoing dynamic interaction with the environment is a conceptual necessity for perceptual experience.

An upshot of this view is that perceptual experience can only supervene on a temporally-extended process, and never on a momentary state. This goes part way to accounting for the experiential quality that Noë observes accompanies presence-as-access, namely the inability to fixate, at an instant, on any atomic visual feature. Note, however, that the content of an experience need not temporally track the material realiser. If it does not, a temporally-extended process could yield the experience of an instant in time in which you take in a shade of red. The extensionalist view that the content does temporally track the vehicle has the benefit of offering an apt way to reformulate the phenomenological account offered by Noë. The inability to take in a shade of red in what you experience as a durationless instant is explained by the fact that having an experience of red is, in part, having an experience of duration. A visual feature never feels phenomenally present at this instant because, before you have a chance to fixate on it, the relevant detail has lapsed from being something that you are confronted with now, to something that you have been confronted with a moment ago. Extensionalism, here, allows that you can experience a shade of red, but stipulates that this takes place – and is experienced as taking place – over an interval of time. The process is likely to be assisted if your movements during that time help you gain information about the patterns of sensorimotor contingency that currently hold between your body and the environment.

Returning now to the debate between Clark and Noë: Clark's objection to presence-as-access is that elapsed portions of the opera singer's note are not accessible now, hence your experience of them cannot be explained by means of their accessibility. Noë's response denies that object and event experience need to be explained in the same way, on the grounds that they are different species of awareness: "objects are primary in our experience [...] experience of events depends on a more basic sensitivity to the presence of objects" (Noë, 2006, p. 31). I earlier argued that Noë's line of response may be inadequate to defend the central tenets of the sensorimotor theory. This drawback can be avoided, I suggest, precisely by reversing Noë's claim, and contending that the experience of objects depends on the ongoing conscious presence of events that have been occurring moments into the past.

7.5 - Chapter Conclusion

It is a matter of some importance that the sensorimotor approach has the resources to provide a plausible answer to the puzzle of temporal experience. Noë (2006) offers an ambivalent response. He suggests, briefly, that it can be explained by temporally-extended ‘coupling’. As I have argued, however, there are no obvious grounds to endorse this, given his main claim that the experience of the opera singer’s note is explained by an understanding, akin to linguistic comprehension, of where the note is coming from. While this latter claim provides an account of temporal experience, it entails an abandonment of a number of the sensorimotor theory’s central tenets, as far as event experience goes. In so doing, it renders the theory more vulnerable to the objection that object perception is not fundamentally sensorimotor either, expressed, for example, by Clark (2008), when he espouses sensorimotor summarising.

With the aim of defending sensorimotor theory, I have aimed to indicate the conceptual foundation of a sensorimotor enactivist account of temporal experience. This involves rejecting the comparison with linguistic understanding and emphasising temporally-extended activity. Adopting extensionalism as a working hypothesis invites, in turn, a particular understanding of object experience. This is worth endorsing in its own right because of the useful gloss it places on Noë’s otherwise compelling ‘virtuality’ metaphor, and the phenomenology of presence-as-access. This consideration gives the sensorimotor theory even more reason to pursue extensionalism as an explanation of temporal experience.

Clark’s (2006) critique has the useful effect of forcing the endorser of sensorimotor theory to evaluate carefully the respective roles of coupling and knowing in the approach. If my suggestion is right, understanding only plays a role if it is directly implicated in bodily coupling with the environment. If bodily interaction is key, then temporal extension is also surely vital. The sensorimotor theory should, therefore, do more than pay lip service to temporality, and take into serious account the temporally-extended nature both of perceptual experience and its material substrate.

8 - Thesis Conclusion

In spite of all the attention the sensorimotor theory has been given since its first official statement, many of its tenets have continued to be incompletely or ambiguously defined, and one of my goals in this thesis has been to go some way toward clarifying them. This undertaking is closely related to the separate project of responding to some of the criticisms that have been levelled against the theory. I have focused in particular on Clark's argument that perceptual experience is likely to depend on the activation of representations which are not intimately involved in bodily movement and which stand-in for types, categories and relative locations rather than laws of sensorimotor contingency. To the end both of responding to Clark's argument and arguing for a particular understanding of the sensorimotor theory, I have suggested that there are good reasons to think that perceptual experience depends on the exercise of behavioural capacities, which may be relatively coarse-grained. In this final passage, I am going to tie together some key strands from the main part of the thesis.

Sensorimotor contingency is perhaps the most important concept deployed by the sensorimotor theory. However, there is sometimes equivocation about what SMCs actually are. One potential tension, although I do not find it particularly problematic, is the shift of emphasis in Noë's sole-author accounts from SMCs governing subdoxastic sense inputs to the SMCs that respectively define P-properties and objective properties, the latter being characterised by the ways P-properties change with movement. More could be said about the relation between these different kinds of property. However, there is nothing obvious to suggest that a final version of the sensorimotor theory could not refer to a combination of all these kinds of sensorimotor contingency.

More problematic is the slide many statements of the sensorimotor account make between describing SMCs as counterfactual conditionals, i.e. the results of merely possible movements, and as patterns of bidirectional output-input dependence manifest in the agent's sensorimotor engagements right now. This distinction is similar to the one suggested by Burhmann, Di Paolo and Barandiaran's concepts of 'sensorimotor environment' and 'sensorimotor habitat'. For the sake of conceptual hygiene, I think SMCs should be

understood only to describe the results of counterfactually possible movements, meaning it need not be possible to identify them from the agent's present bodily interactions. SMCs, on this understanding, describe the properties a subject exhibits an appropriate kind of sensitivity to if she perceives, but do not presuppose that the subject does perceive. The concept, so construed, avoids prejudging the means by which subjects exhibit the appropriate sensitivity, which could involve the activation of representations, the exercise of behavioural capacities, or both.

Further, adopting the counterfactual conditional understanding of SMCs does justice to the sensorimotor theory's central suggestion that perceivers, by virtue of possessing sensorimotor mastery, experience the presence (in absence) of parts of the environment even when they are not presently interacting with those features in an outward bodily way. I have sometimes heard people suggest that SMCs are properties a perceiver can 'have'. This is inaccurate. They could be properties a perceiver's interactions with the environment have. On the understanding I prefer, they are bodily relations a perceiver stands in to the outside environment which we describe by specifying what would happen *if* the subject or objects in question moved in particular ways. One reason it is important to clearly understand the concept of sensorimotor contingency is that a next phase for the sensorimotor theory will be to give detailed descriptions of the SMCs that comprise specific experiential properties, such as shades of red. Moreover, how one understands the concept of sensorimotor contingency will have a knock on effect on how one understands the concept of sensorimotor mastery as well as the physical mechanism implicated in it.

As I suggested earlier in the thesis, sensorimotor mastery admits two possible directions of fit. It is sometimes glossed as Being Able To Act, which implies a desire-like direction of fit, and sometimes as Knowing What Would Happen If You Did Act, which implies a belief-like direction of fit. One proposal is that sensorimotor mastery is essentially belief-like, and that having an ability to act is merely a trivial consequence of possessing sensorimotor mastery. Noë offers this interpretation at times, although not consistently. The construal of sensorimotor mastery in which it is primarily belief-like invites the suggestion that the exercise of sensorimotor mastery is constituted by the deployment of internal representations bearing content about SMCs. The representations could take the form, for

instance, of the predictive models proposed by Seth (2014). Although this sort of account could turn out to be correct, I do not think endorsing it is the sensorimotor theorist's best bet, since the proposal entails abandoning many of the improvements the sensorimotor theory otherwise makes on the orthodox approach.

One improvement the sensorimotor theory would otherwise make is to avoid placing an unsustainable explanatory burden on the concept of internal representation. There are significant limitations to the explanatory efficacy of appeals to internal representation. Hutto and Myin (2012) have demonstrated that it is not clear that subpersonal content can ever be accounted for in a naturalistically respectable way. Even if one were not moved by Hutto and Myin's contention that any and all content-talk should therefore be avoided, there are sensible limits that should be placed on internal representation's theoretical role. Since it has not been established that neural states can bear metaphysically real content rather than content we ascribe as a useful metaphor because it is explanatorily expedient to do so, we should not appeal to content to account for phenomenal qualities, particularly if we, as physicalists, hope to respond to those with metaphysically-motivated anti-physicalist intuitions about qualia. Setting this point to one side, it is not certain, in any case, that appealing solely to the content and functional role of neural representations will ever be sufficient by itself to make it intelligible how personal level perceptual content or perceptual phenomenology is connected to the physical.

These drawbacks also speak against the objection, from Clark (2008), that perceptual experience depends on the activation in the ventral stream of internal representations geared towards the exercise of epistemic skills like sorting, sifting and classifying, and that the sensorimotor theory is therefore implausible. If these skills are not bodily skills, the view ends up depending on the content of the internal representations to account for phenomenal qualities and, as I have suggested, internal representations do not appear to be up to this job. If sorting and sifting are varieties of sensorimotor skill, as is plausible, then Clark's account of ventral stream perception does not challenge the sensorimotor theory's core claim that perception is constituted by the exercise of sensorimotor skills, which we can allow are likely to be coarse-grained rather than fine-grained.

Sensorimotor theorists also have a strategic reason not to endorse the claim that the

exercise of sensorimotor understanding is constituted by the activation of representations, since doing so makes the theory a hostage to fortune. If the phenomenology of presence-in-absence and the logical relation between perception and action could be compellingly accounted for in a way that no longer entails that perceivers master SMCs, then few of the sensorimotor theory's core tenets would remain secure. Having said that, if it were established the perceptual experience does depend on nothing more than the activation of internal representations, the sensorimotor theory would continue to have promise as an account of phenomenal qualities. Internal representations, supposing that they do bear genuine content, could not bear the hyperintensional content required to actively discount the possibility that representations of types, categories and locations are also representations of SMCs. Appealing to SMCs rather than types, categories and locations could, however, help representationalists about phenomenal qualities respond to the Inverted Earth counterexample.

Although there are problems with tying the sensorimotor theory too heavily to the concept of internal representation, sensorimotor mastery must, at least in a loose sense, involve Knowing What Would Happen If You Did Act, since SMCs describe the consequences of merely possible movements, and perceivers must grasp them in some way. My suggestion is that sensorimotor mastery nonetheless does not depend essentially on internal representation, but on behavioural capacities. These capacities, I suggest, implicate Knowing What Would Happen If You Did Act just in case they meet appropriate criteria. This approach perhaps comes close to what Noë would say is the behaviourist mistake of conflating effect (i.e. behaviour) with cause (i.e. knowledge). However, it avoids crude behaviourism. First, the idea is not simply to abandon sensorimotor mastery as a belief-like capacity and appeal to behavioural capacities instead. My suggestion is that the relevant behavioural capacities logically implicate the appropriate understanding, meaning that the understanding is a real entity even if it is not reified as something like a neural representation. Thus it does not deny perceptual consciousness, but explains perceptual consciousness by appeal to sensorimotor mastery. Further, the present view avoids the implausible suggestion that perceivers must in every case move their bodies in order to perceive, since the behavioural capacities may be active merely by virtue of a subject's being *prepared* to behave. Preparedness

to behave could, in principle, be ascribed on the basis of neural and bodily states without appeal to representations or presently-occurring bodily movements.

Finally, the view I am recommending does not rule out the possibility of giving a restricted theoretical role to internal representation. In fact, it endorses pluralism about how the relevant behavioural capacities may be accounted for physicalistically. It is compatible with an entirely nonrepresentational dynamical systems approach. It is also compatible with an approach that appeals to information processing, including the deployment of internal representations, should they prove the most effective way to explain how the appropriate behavioural capacities are realised. This approach to representation avoids unduly conflating the relatively innocent suggestion that representation is a useful metaphorical posit for making sense of the physical processes implicated in perception with the faulty suggestion that perception just is the activation of an internal representation.

A final thing is to see how the behavioural capacity view can make sense of the phenomenology of object experience described by Noë, as well as provide a distinctive answer to the puzzle of temporal experience, which present a specific challenge to the sensorimotor theory as well as presenting a challenge to accounts of perception more generally. This is my suggestion: Perception unfolds as a temporally-extended process involving the exercise of capacities for bodily behaviour. This process frequently involves outwardly observable bodily movement, although it need not, since behavioural capacities are often exercised merely in the course of becoming prepared to move our bodies. This process cannot be broken down, in a principled way, into discrete time slices. Considered at any particular instant, however, a perceiver is poised to make certain bodily movements. This poise reflects the perceiver's understanding of certain SMCs, where SMCs are taken to be characterised by counterfactual conditionals specifying the results of possible movements. The movements a perceiver is poised to make at a particular instant are sufficient to account for the perceiver's experiencing the presence in absence of certain environmental detail. However the SMCs they are exercising mastery of at this instant are insufficient to determinately specify any particular environmental property, for example a shade of red.

The experience of a determinate feature, such as a shade of red, only happens in the course of the temporally-extended exercise of that behavioural capacity. Over a period of time,

a perceiver undergoes a continually evolving poise to make skillful bodily movements. Over the course of this evolution, the perceiver may experience a shade of red, although they can never fixate on the shade of red, since seeing the shade of red depends on this evolving, temporally extended, non-atomic process. If it the content of the experience did not temporally track the vehicle, there need not be any obstacle to seeing the shade of red all at once, as a non-temporally extended content could supervene on a temporally-extended vehicle. But we can make sense of Noë's description of visual phenomenology, which I find compelling, by supposing that the content temporally tracks the vehicle. This means the experience of the shade of red is an experience of temporal extension. This is suggestive of a different response to the opera singer case than the one given by Noë, namely a version of extensionalism rather than the cinematic approach he endorses. This gives us a way to account for the phenomenology associated both with event and object experience.

One thing I have not addressed in detail in this thesis is the precise mechanism implicated in sensorimotor mastery, and this aspect of the sensorimotor theory continues to be in need of significant development. Further, I have not examined in detail the precise character of the behavioural capacities that would implicate different forms of sensorimotor mastery; offering such an account has the potential to be a useful next step for the sensorimotor theory. A final noteworthy limitation is that I have not resolved the more general question of how to licence the claim that appropriate behaviours and behavioural capacities are themselves intentionally-directed. Answering this question might involve endorsing the biological enactivists' appeal to autonomy and normativity, although we might finally decide that notions like these are dispensable. This is a question that the sensorimotor theory ought to take a stance on. However, I believe that with a clear understanding of the fundamentals of the sensorimotor theory, and an appreciation of its significant merits, we ought to be well placed to progress the theory further.

Bibliography

- Aizawa, K. (2010). Consciousness: Don't give up on the brain. *Royal Institute of Philosophy Supplements*, 67, 263–284.
- Allen, C. (2009). Teleological Notions in Biology. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2009).
- Ballard, D. H., Hayhoe, M. M., Pook, P. K., & Rao, R. P. N. (1997). Deictic codes for the embodiment of cognition. *Behavioral and Brain Sciences*, 20(4), 723–742.
- Bennett, M. R., & Hacker, P. M. S. (2003). *Philosophical Foundations of Neuroscience*. Malden, Mass.: Wiley-Blackwell.
- Berlin, B., & Kay, P. (1969). *Basic Color Terms: their Universality and Evolution*. Berkeley and Los Angeles: University of California Press.
- Bermudez, J. L. (2000). Personal and subpersonal: A difference without a distinction. *Philosophical Explorations*, 3(1), 63–82.
- Blackmore, S. J. (2002). There is no stream of consciousness. *Is the Visual World a Grand Illusion? Spec. issue of Journal of Consciousness Studies*, 9(5-6), 17–28.
- Block, N. (1990a). Consciousness and accessibility. *Behavioral and Brain Sciences*, 13(4), 596–598.
- Block, N. (1990b). Inverted Earth. *Philosophical Perspectives*, 4, 53–79.
- Block, N. (1996). Mental paint and mental latex. *Philosophical Issues*, 7, 19.
- Block, N. (1998). Is experiencing just representing? *Philosophy and Phenomenological Research*, 58(3), 663–670.
- Block, N. (2001). Behaviorism revisited. *Behavioral and Brain Sciences*, 24(05), 977–978.
- Block, N. (2005). Review of Alva Noë, Action in Perception. *Journal of Philosophy*, 102(5), 259–272.
- Brooks, R. (1991). Intelligence without representation. *Artificial Intelligence*, 47, 139–159.

- Buhrmann, T., Di Paolo, E. A., & Barandiaran, X. (2013). A dynamical systems account of sensorimotor contingencies. *Frontiers in Psychology*, 4.
- Burge, T. (1979). Individualism and the mental. *Midwest Studies in Philosophy*, 4(1), 73–122.
- Campbell, J. (2002). *Reference and Consciousness*. Oxford: Oxford University Press.
- Campbell, J. (2008). Sensorimotor knowledge and naïve realism. *Philosophy and Phenomenological Research*, 76(3), 666–673.
- Carruthers, P. (2011). Higher-order theories of consciousness. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2011.).
- Chalmers, D. J. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. Oxford: Oxford University Press.
- Chemero, A. (2009). *Radical Embodied Cognitive Science*. Cambridge, Mass.: The MIT Press.
- Chemero, A., & Cordeiro, W. (2000). Dynamical, Ecological Sub-Persons. *A Field Guide to the Philosophy of Mind*.
- Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Ryu, S. I., & Shenoy, K. V. (2010). Cortical preparatory activity: Representation of movement or first cog in a dynamical machine? *Neuron*, 68(3), 387–400.
- Clark, A. (1997). *Being There: Putting Brain, Body, and World Together Again*. Cambridge, Mass.: The MIT Press.
- Clark, A. (2006). That lonesome whistle: A puzzle for the sensorimotor model of perceptual experience. *Analysis*, 66(289), 22–25.
- Clark, A. (2008). *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford: Oxford University Press.
- Clark, A. (2009). Spreading the joy? Why the machinery of consciousness is (probably) still in the head. *Mind*, 118(472), 963–993.
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(3), 181–204.
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58(1), 7–19.

- Clark, A., & Toribio, J. (2001). Sensorimotor chauvinism? *Behavioral and Brain Sciences*, 24(05), 979–980.
- Crick, F., & Koch, C. (2003). A framework for consciousness. *Nature Neuroscience*, 6(2), 119–126.
- Dainton, B. (2000). *Stream of Consciousness*. London: Routledge.
- Dainton, B. (2010). Temporal Consciousness. *The Stanford Encyclopedia of Philosophy, Fall 2010 Edition*. Internet Encyclopedia.
- Davidson, D. (1987). Knowing one's own mind. *Proceedings and Addresses of the American Philosophical Association*, 60(3), 441–458.
- Degenaar, J., & Myin, E. (2013). The structure of color experience and the existence of surface colors. *Philosophical Psychology*, (3), 1–17.
- Dennett, D. C. (1978). *Brainstorms: Philosophical Essays on Mind and Psychology*. Cambridge, Mass.: The MIT Press.
- Dennett, D. C. (1987). *The Intentional Stance*. Cambridge, Mass.: The MIT Press.
- Dennett, D. C. (1991). *Consciousness Explained*. Boston: Little, Brown and Co.
- Dennett, D. C. (1992). "Filling in" versus finding out: A ubiquitous confusion in cognitive science. In P. Van de Broek & D. C. Knill (Eds.), *Cognition: Conceptual and Methodological Issues* (pp. 33–49). Washington, DC: American Psychological Association.
- Dennett, D. C. (2001). Surprise, surprise. *Behavioral and Brain Sciences*, 24(05), 982–982.
- Dennett, D. C. (2002). How could I be wrong? How wrong could I be? *Is the Visual World a Grand Illusion? Spec. issue of Journal of Consciousness Studies*, 9(5-6), 17–28.
- Dennett, D. C. (2007). Philosophy as naive anthropology: Reply to Bennett and Hacker. In *Neuroscience and Philosophy: Brain, Mind, and Language* (1 edition.). Columbia University Press.
- Descartes, R. (1985b). Meditations on First Philosophy. In J. Cottingham, R. Stoothoff, & D. Murdoch (Trans.), *The Philosophical Writings of Descartes: Volume 2*. Cambridge, UK: Cambridge University Press.

- Descartes, R. (1985a). Optics. In J. Cottingham, R. Stoothoff, & D. Murdoch (Trans.), *The Philosophical Writings of Descartes: Volume 1*. Cambridge, UK: Cambridge University Press.
- Drayson, Z. (2014). The personal/subpersonal distinction. *Philosophy Compass*, 9(5), 338–346.
- Dretske, F. (1981). *Knowledge and the Flow of Information*. Cambridge, Mass.: The MIT Press.
- Dretske, F. (1995). *Naturalizing the Mind*. Cambridge, Mass.: The MIT Press.
- Dreyfus, H. L. (2013). The myth of the pervasiveness of the mental. In J. K. Schear (Ed.), *Mind, Reason and Being-in-the-World* (pp. 15–40). London: Routledge.
- Evans, G. (1982). *The Varieties of Reference*. Oxford: Oxford University Press.
- Evans, G. (1985). Molyneux’s question. In *Collected Papers*. Oxford: Oxford University Press.
- Fodor, J. A. (1983). *The Modularity of Mind: An Essay on Faculty Psychology*. Cambridge, Mass.: The MIT Press.
- Fodor, J. A. (1990). A theory of content I: The problem. In *A Theory of Content and Other Essays* (pp. 51–88). Cambridge, Mass.: The MIT Press.
- Gangopadhyay, N., & Kiverstein, J. (2009). Enactivism and the unity of perception and action. *Topoi*, 28(1), 63–73.
- Geach, P. T. (1969). *God and the Soul*. London: Routledge.
- Geldard, F. A., & Sherrick, C. E. (1972). The cutaneous “rabbit”: A perceptual illusion. *Science*, 178(4057), 178–179.
- Gendler, T. S., & Hawthorne, J. (2002). *Conceivability and Possibility*. Oxford: Oxford University Press.
- Glock, H.-J. (2008). Animal minds: Conceptual problems. Presented at the Inaugural Lecture, University of Zurich.
- Goldman, A., & de Vignemont, F. (2009). Is social cognition embodied? *Trends in cognitive sciences*, 13(4), 154–159.
- Goodman, N. (1978). *Ways of Worldmaking*. Brighton: Harvester.

- Grush, R. (2007). Time and experience. In *The Philosophy of Time*. Frankfurt: Klosterman.
- Harvey, I., Husbands, P., & Cliff, D. (1994). Seeing the light: Artificial evolution, real vision. In D. Cliff, P. Husbands, J.-A. Meyer, & S. W. Wilson (Eds.), *SAB94 Proceedings of the Third International Conference on Simulation of Adaptive Behavior: From Animals to Animats 3*. Cambridge, Mass.: The MIT Press.
- Hohwy, J. (2013). *The Predictive Mind*. Oxford: Oxford University Press.
- Hurley, S. (1998). *Consciousness in Action*. Cambridge, Mass.: Harvard University Press.
- Hurley, S., & Noë, A. (2003). Neural plasticity and consciousness. *Biology and Philosophy*, 18, 131–168.
- Hurley, S., & Noë, A. (2007). Can Hunter-Gatherers Hear Color? In M. Smith, R. Goodin, & G. Geoffrey (Eds.), *Common Minds: Essays in Honor of Philip Pettit* (pp. 55–83). Oxford: Oxford University Press.
- Husserl, E. (1991). *On the Phenomenology of the Consciousness of Internal Time (1893 - 1917)*. (J. B. Brough, Trans.). Dordrecht: Springer.
- Hutto, D. D., Kirchhoff, M., & Myin, E. (2014). *Extensive enactivism: Why keep it all in?* Unpublished manuscript.
- Hutto, D. D., & Myin, E. (2012). *Radicalizing Enactivism: Basic Minds Without Content*. Cambridge, Mass.: The MIT Press.
- Jackson, F. (1982). Epiphenomenal qualia. *Philosophical Quarterly*, 32(127), 127–136.
- James, W. (1981). *The Principles of Psychology, Volumes I and II*. Cambridge, Mass.: Harvard University Press.
- Kant, I. (1998). *Critique of Pure Reason*. (P. Guyer & A. W. Wood, Trans.). Cambridge, UK: Cambridge University Press.
- Kelly, S. (2005). The Puzzle of Temporal Experience. In *Cognition and the Brain: The Philosophy and Neuroscience Movement*. Cambridge, UK: Cambridge University Press.
- Kelly, S. D. (2008). Content and constancy: Phenomenology, psychology, and the content of perception. *Philosophy and Phenomenological Research*, 76(3), 682–690.

- Kenny, A. J. P. (1971). The homunculus fallacy. In M. Grene (Ed.), *Interpretations of Life and Mind* (pp. 65–74). London: Routledge.
- Khalidi, M. A. (2007). Innate cognitive capacities. *Mind & Language*, 22(1), 92–115.
- Kohler, I. (1962). Experiments with goggles. *Scientific American*, 206, 62–72.
- Kosslyn, S. M. (1994). *Image and Brain: The Resolution of the Imagery Debate*. Cambridge, Mass.: The MIT Press.
- Lehar, S. (1998). Gestalt isomorphism and the primacy of the subjective perceptual experience. *Behavioral and Brain Sciences*, 21(06), 763–764.
- Lehar, S. (2003). Gestalt isomorphism and the primacy of subjective conscious experience: A gestalt bubble model. *Behavioral and Brain Sciences*, 26(04), 375–408.
- Lewis, D. (1983). Postscript to “Mad pain and Martian pain.” In *Philosophical Papers: Volume 1* (pp. 130–132). Oxford: Oxford University Press.
- Lindberg, D. C. (1976). *Theories of Vision from Al-Kindi to Kepler*. Chicago: University of Chicago Press.
- Llinás, R., & Ribary, U. (1994). Coherent 40-Hz oscillation characterizes dream state in humans. *Proceedings of the National Academy of Sciences of the United States of America*, 90(5), 2078–2081.
- Lycan, W. G. (1987). *Consciousness*. Cambridge, Mass.: The MIT Press.
- Marr, D. (1982). *Vision*. Cambridge, Mass.: The MIT Press.
- Martin, M. (2008). Commentary on Action in Perception. *Philosophy and Phenomenological Research*, 76(3), 674–681.
- Matin, E. (1974). Saccadic suppression: A review and an analysis. *Psychological Bulletin*, 81(12), 899–917.
- Maturana, H., & Varela, F. (1980). *Autopoiesis and Cognition: The Realization of the Living*. Dordrecht: Springer.
- McConkie, G. W., & Zola, D. (1979). Is visual information integrated across successive fixations in reading? *Perception and Psychophysics*, 25, 221–24.

- McDowell, J. (1994a). *Mind and World*. Cambridge, Mass.: Harvard University Press.
- McDowell, J. (1994b). The content of perceptual experience. *The Philosophical Quarterly*, 44(175), 190–205.
- McGinn, C. (1991). *The Problem of Consciousness: Essays Toward a Resolution*. Oxford: Blackwell.
- Merleau-Ponty, M. (2013). *Phenomenology of Perception*. (D. A. Landes, Trans.). London: Routledge.
- Millikan, R. G. (1984). *Language, Thought, and Other Biological Categories: New Foundations for Realism*. Cambridge, Mass.: The MIT Press.
- Millikan, R. G. (1993). Compare and Contrast Dretske, Fodor, and Millikan on Teleosemantics. In *White Queen Psychology and Other Essays for Alice* (pp. 123–134). Cambridge, Mass.: The MIT Press.
- Milner, D., & Goodale, M. A. (1995). *The Visual Brain in Action*. Oxford: Oxford University Press.
- Montero, B. (1999). The body problem. *Noûs*, 33(2), 183–200.
- Müller, J. (1838). *Handbuch der Physiologie des Menschen für Vorlesungen*. Coblenz: J. Hölscher.
- Myin, E. (2001). Constrained inversions of sensations. *Philosophica*, 68(2), 31–40.
- Myin, E., & Degenaar, J. (2014). Enactive vision. In L. Shapiro (Ed.), *The Routledge Handbook of Embodied Cognition* (pp. 90–98). London: Routledge.
- Myin, E., & O'Regan, J. K. (2002). Perceptual consciousness, access to modality and skill theories: A way to naturalize phenomenology? *Journal of Consciousness Studies*, 9(1), 27–45.
- Nakayama, K. (1994). James J. Gibson - an appreciation. *Psychological Review*, 101(2), 329–335.
- Noë, A. (2002). Is the Visual World a Grand Illusion? *Is the Visual World a Grand Illusion? Special issue of Journal of Consciousness Studies*, 9(5-6), 1–12.
- Noë, A. (2004). *Action in Perception*. Cambridge, Mass.: The MIT Press.
- Noë, A. (2006). Experience of the world in time. *Analysis*, 66(289), 26–32.

- Noë, A. (2008). Reply to Campbell, Martin, and Kelly. *Philosophy and Phenomenological Research*, 76(3), 691–706.
- Noë, A. (2010). Vision without representation. In N. Gangopadhyay, M. Madary, & F. Spicer (Eds.), *Perception, action, and consciousness: Sensorimotor Dynamics and Two Visual Systems* (pp. 245–256). Oxford: Oxford University Press.
- Noë, A. (2012). *Varieties of Presence*. Cambridge, Mass.: Harvard University Press.
- Noë, A., & Thompson, E. (2004). Are there neural correlates of consciousness? *Journal of Consciousness studies*, 11(1), 3–28.
- O'Brien, G., & Opie, J. (2009). The role of representation in computation. *Cognitive Processing*, 10(1), 53–62.
- O'Regan, J. K. (1984). Retinal versus extraretinal influences in flash localization during saccadic eye movements in the presence of a visible background. *Perception & Psychophysics*, 36(1), 1–14.
- O'Regan, J. K. (1992). Solving the “real” mysteries of visual perception: The world as an outside memory. *Canadian Journal of Psychology*, 46(3), 461.
- O'Regan, J. K. (2011). *Why Red Doesn't Sound Like a Bell: Understanding the Feel of Consciousness*. Oxford: Oxford University Press.
- O'Regan, J. K., & Block, N. (2012). Discussion of J. Kevin O'Regan's “Why red doesn't sound like a bell: Understanding the feel of consciousness.” *Review of Philosophy and Psychology*, 3(1), 89–108.
- O'Regan, J. K., Myin, E., & Noë, A. (2005). Skill, corporality and alerting capacity in an account of sensory consciousness. *Progress in Brain Research*, 150, 55–68.
- O'Regan, J. K., Myin, E., & Noë, A. (2004). Towards an analytic phenomenology: The concepts of “bodiliness” and “grabbiness.” In A. Carsetti (Ed.), *Seeing, Thinking and Knowing* (Vol. 38, pp. 103–114). Springer Netherlands.
- O'Regan, J. K., & Noë, A. (2001a). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24(5), 883–917.
- O'Regan, J. K., & Noë, A. (2001b). Acting out our sensory experience. *Behavioral and Brain Sciences*, 24(5), 1011–1030.

- O'Regan, J. K., Rensink, R. A., & Clark, J. J. (1999). Change-blindness as a result of “mudsplashes.” *Nature*, *398*(6722), 34–34.
- Penrose, R. (1994). *Shadows of the Mind: A Search for the Missing Science of Consciousness*. Oxford: Oxford University Press.
- Pessoa, L., Thompson, E., & Noë, A. (1998). Finding out about filling-in: A guide to perceptual completion for visual science and the philosophy of perception. *Behavioral and Brain Sciences*, *21*(6), 723–802.
- Pettit, P. (2003). Looks as powers. *Philosophical Issues*, *13*(1), 221–252.
- Philipona, D. L., & O'Regan, J. K. (2006). Color naming, unique hues, and hue cancellation predicted from singularities in reflection properties. *Visual neuroscience*, *23*(3-4), 331–339.
- Port, R. F., & Gelder, T. V. (1995). *Mind as Motion: Explorations in the Dynamics of Cognition*. Cambridge, Mass.: The MIT Press.
- Prinz, J. J. (2006). Putting the brakes on enactive perception. *Psyche*, *12*(1), 1–19.
- Putnam, H. (1975). The meaning of “meaning.” *Minnesota Studies in the Philosophy of Science*, *7*, 131–193.
- Ramachandran, V. S., & Gregory, R. L. (1991). Perceptual filling in of artificially induced scotomas in human vision. *Nature*, *350*(6320), 699–702.
- Ramsey, W. M. (2007). *Representation Reconsidered*. Cambridge, UK: Cambridge University Press.
- Rensink, R. A., O'Regan, J. K., & Clark, J. J. (1997). To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science*, *8*(5), 368–373.
- Roberts, T. (2010). Understanding “sensorimotor understanding.” *Phenomenology and the Cognitive Sciences*, *9*(1), 101–111.
- Rorty, A. O. (1971). Not every homunculus spoils the argument. In M. Grene (Ed.), *Interpretations of Life and Mind* (pp. 75–80). London: Routledge.
- Rowlands, M. (2006). *Body Language: Representation in Action*. Cambridge, Mass.: The MIT Press.

- Rowlands, M. (2010). *The New Science of the Mind: From Extended Mind to Embodied Phenomenology*. Cambridge, Mass.: The MIT Press.
- Rupert, R. D. (2011). Embodiment, consciousness, and the massively representational mind. *Philosophical Topics*, 39(1), 99–120.
- Ryle, G. (1949). *The Concept of Mind*. London: Hutchinson's University Library.
- Ryle, G. (1979). *On Thinking*. (K. Kolenda, Ed.). Oxford: Blackwell.
- Searle, J. R. (1994). Animal Minds. *Midwest Studies In Philosophy*, 19(1), 206–219.
- Seth, A. K. (2014). A predictive processing theory of sensorimotor contingencies: Explaining the puzzle of perceptual presence and its absence in synesthesia. *Cognitive Neuroscience*, 5(2), 97–118.
- Shoemaker, S. (1980). Causality and properties. In P. Van Inwagen (Ed.), *Time and Cause* (pp. 109–135). Dordrecht: Reidel.
- Silberstein, M., & Chemero, A. (2012). Complexity and extended phenomenological-cognitive systems. *Topics in Cognitive Science*, 4(1), 35–50.
- Simons, D. J., & Chabris, C. F. (1999). Gorillas in our midst: Sustained inattention blindness for dynamic events. *Perception*, 28(9), 1059–1074.
- Sprevak, M. (2013). Fictionalism about neural representations: *The Monist*, 96(4), 539–560.
- Stanley, J., & Williamson, T. (2001). Knowing how. *Journal of Philosophy*, 98(8), 411–444.
- Summerfield, C., Trittschuh, E. H., Monti, J. M., Mesulam, M.-M., & Egner, T. (2008). Neural repetition suppression reflects fulfilled perceptual expectations. *Nature Neuroscience*, 11(9), 1004–1006.
- Tanji, J., & Evarts, E. V. (1976). Anticipatory activity of motor cortex neurons in relation to direction of an intended movement. *Journal of Neurophysiology*, 39(5), 1062–1068.
- Teller, D. Y., & Pugh, E. N. (1983). Linking propositions in color vision. In J. D. Mollon & L. T. Sharpe (Eds.), *Colour Vision: Physiology and Psychophysics*. London: Academic Press.
- Thelen, E., & Smith, L. B. (1996). *A Dynamic Systems Approach to the Development of Cognition and Action*. Cambridge, Mass.: The MIT Press.

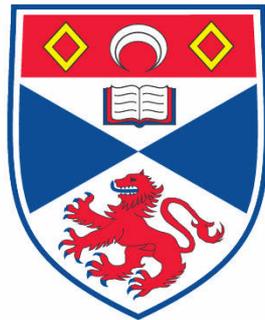
- Thompson, E. (2007). *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*. Cambridge, Mass.: Harvard University Press.
- Tye, M. (2000). *Consciousness, Color and Content*. Cambridge, Mass.: The MIT Press.
- Van Gelder, T. (1998). The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences*, 21(5), 615–628.
- Varela, F. (1999). The Specious Present: A Neurophenomenology of Time Consciousness. In J. Petitot, F. J. Varela, B. Pacoud, & J.-M. Roy (Eds.), *Naturalizing Phenomenology* (pp. 266–314). Stanford: Stanford University Press.
- Varela, F., Thompson, E., & Rosch, E. (1991). *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge, Mass.: The MIT Press.
- Ward, D. (2009). The agent in magenta. *Psyche*, 15(1).
- Ward, D., Roberts, T., & Clark, A. (2011). Knowing what we can do: Actions, intentions, and the construction of phenomenal experience. *Synthese*, 181(3), 375–394.
- Ward, D., & Stapleton, M. (2012). Es are good: Cognition as enacted, embodied, embedded, affective and extended. In F. Paglieri (Ed.), *Consciousness in Interaction: The role of the natural and social context in shaping consciousness* (Vol. 86). Amsterdam: John Benjamins.
- Wheeler, M. (2005). *Reconstructing the Cognitive World: The Next Step*. Cambridge, Mass.: The MIT Press.
- Wheeler, M. (2010). In defence of extended functionalism. In R. Menary (Ed.), *The Extended Mind*. Cambridge, Mass.: The MIT Press.
- Wheeler, M. (2013). Science friction: Phenomenology, naturalism and cognitive science. *Royal Institute of Philosophy Supplements*, 72, 135–167.
- Wheeler, M. (2014). *The revolution will not be optimised: Enactivism, embodiment and relationality*. Presented at the 50th Anniversary Convention of the AISB, London.
- Wheeler, M., & Clark, A. (1999). Genic representation: Reconciling content and causal complexity. *British Journal for the Philosophy of Science*, 50(1), 103–135.
- Whitham, E. M., Fitzgibbon, S. P., Lewis, T. W., Pope, K. J., DeLosAngeles, D., Clark, C. R., ... Willoughby, J. O. (2011). Visual experiences during paralysis. *Frontiers in Human Neuroscience*, 5.

Wittgenstein, L. (2009). *Philosophical Investigations*. (G. E. M. Anscombe, J. Schulte, & P. M. S. Hacker, Trans.) (4th Edition edition.). Malden, Mass.: Wiley-Blackwell.

Wolfe, J. M. (1997). Visual experience: Less than you think, more than you remember. In C. Taddei-Ferretti & C. Musio (Eds.), *Neuronal Bases and Psychological Aspects of Consciousness: Proceedings of the International School of Biocybernetics* (pp. 165–185). Singapore: World Scientific.

The Sensorimotor Theory of Perceptual Experience

David Silverman



This thesis is submitted in partial fulfilment for the degree of PhD at the University of St Andrews and University of Stirling Graduate Programme in Philosophy (SASP)

30th June 2014

Abstract

The sensorimotor theory is an influential, non-mainstream account of perception and perceptual consciousness intended to improve in various ways on orthodox theories. It is often taken to be a variety of enactivism, and in common with enactivist cognitive science more generally, it de-emphasises the theoretical role played by internal representation and other purely neural processes, giving theoretical pride of place instead to interactive engagements between the brain, non-neural body and outside environment.

In addition to offering a distinctive account of the processing that underlies perceptual consciousness, the sensorimotor theory aims to offer a new and improved account the logical and phenomenological character of perceptual experience, and the relation between physical and phenomenal states. Since its inception in a 2001 paper by O'Regan and Noë, the theory has prompted a good deal of increasingly prominent theoretical and practical work in cognitive science, as well as a large body of secondary literature in philosophy of cognitive science and philosophy of perception. In spite of its influential character, many of the theory's most basic tenets are incompletely or ambiguously defined, and it has attracted a number of prominent objections.

This thesis aims to clarify the conceptual foundations of the sensorimotor theory, including the key theoretical concepts of sensorimotor contingency, sensorimotor mastery, and presence-as-access, and defends a particular understanding of the respective theoretical roles of internal representation and behavioural capacities. In so doing, the thesis aims to highlight the sensorimotor theory's virtues and defend it from some leading criticisms, with particular attention to a response by Clark which claims that perception and perceptual experience plausibly depend on the activation of representations which are not intimately involved in bodily engagements between the agent and environment. A final part of the thesis offers a sensorimotor account of the experience of temporally extended events, and shows how with reference to this we can better understand object experience.

1. Candidate's declarations:

I, David Silverman, hereby certify that this thesis, which is approximately 75,000 words in length, has been written by me, and that it is the record of work carried out by me, or principally by myself in collaboration with others as acknowledged, and that it has not been submitted in any previous application for a higher degree.

I was admitted as a research student in September 2012 and as a candidate for the degree of PhD in September 2010; the higher study for which this is a record was carried out in the University of St Andrews between 2010 and 2014.

Date signature of candidate

2. Supervisor's declaration:

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree.

Date signature of supervisor

3. Permission for publication: (to be signed by both candidate and supervisor)

In submitting this thesis to the University of St Andrews I understand that I am giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. I also understand that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker, that my thesis will be electronically accessible for personal or research use unless exempt by award of an embargo as requested below, and that the library has the right to migrate my thesis into new electronic forms as required to ensure continued access to the thesis. I have obtained any third-party copyright permissions that may be required in order to allow such access and migration, or have requested the appropriate embargo below.

The following is an agreed request by candidate and supervisor regarding the publication of this thesis:

PRINTED COPY

a) No embargo on print copy

ELECTRONIC COPY

a) No embargo on electronic copy

Date signature of candidate signature of supervisor.....

Contents

Acknowledgements.....	v
1 - Introduction: The Sensorimotor Theory of Perceptual Experience.....	1
1.1 - 4E Cognitive Science and the Varieties of Enactivism.....	3
1.2 - Structure of Thesis.....	6
2 - Making Conceptual Space.....	8
2.1 - Kepler and Descartes: The Lessons from History.....	8
2.1.1 - Pictures, Filling-In and the Bridge Locus.....	10
2.1.2 - Neural Representation.....	16
2.1.3 - Varieties of Homuncularity.....	19
2.2 - Conservative and Radical Strands.....	23
2.2.1 - The Tenets of Sensorimotor Theory: A Brief Outline.....	25
3 - The Tenets of Sensorimotor Theory.....	31
3.1 - The World as an Outside Memory.....	31
3.2 - Phenomenal Qualities: The Sensorimotor Contingency Account.....	37
3.2.1 - The "Qualia" Problem.....	38
3.2.2 - Qualia and Neural-Correlates of Consciousness.....	39
3.2.3 - The Sensorimotor Contingency Account.....	41
3.2.4 - SMCs as Counterfactual Conditionals.....	44
3.2.5 - Perspectival Properties and Objective Properties.....	47
3.2.6 - Bodiliness, Insubordinateness, Richness and Grabbiness.....	51
3.3 - Sensorimotor Mastery.....	53
3.3.1 - Sensorimotor Attunement.....	54
3.3.2 - Sensorimotor Tracking.....	62
3.3.3 - Sensorimotor Understanding.....	65
3.4 - Perceptual Presence: As-Access and In-Absence.....	67
3.5 - Perception and Dynamical Entanglement.....	77
3.6 - The Extensive Conscious Mind.....	81

4 - Mastering Sensorimotor Understanding.....	93
4.1 - 'Action in Perception' in Action in Perception.....	93
4.2 - Which Direction of Fit?.....	99
4.3 - Sensorimotor Understanding as Subpersonal Representation.....	104
4.4 - Sensorimotor Understanding in Bodily Capacities.....	109
5 - Phenomenal Character: A Sensorimotor-Representation Hybrid.....	116
5.1 - The Non-Sensorimotor, Skill-Based Approach.....	116
5.2 - The Policy of Containment: A Conservative Response.....	122
5.3 - A Challenge to the Policy of Containment.....	124
5.3.1 - The WAY problem.....	125
5.3.2 - Sense and Sensorimotor Contingencies.....	128
5.4 - Mary and Inverted Earth: A Sensorimotor Response.....	134
5.5 - Relation to Rupert's 'Massively Representational Mind'.....	139
5.6 - Chapter Conclusion.....	143
6 - Sorting, Sifting and Other Sensorimotor Skills.....	144
6.1 - Avoiding Zombie Intuitions.....	144
6.2 - Sorting and Sifting as Sensorimotor Skills.....	149
6.3 - Homunculi that Sort and Sift.....	153
6.4 - Radical Anti-Representationalism.....	157
6.5 - Representation Restricted.....	163
7 - Object Experience as Temporal Experience.....	170
7.1 - The Puzzle of Temporal Experience.....	170
7.2 - Coupling and Knowing in Temporal Experience.....	174
7.3 - Object Experience and Skill-Based Access.....	176
7.4 - Extensionalism About Event and Object Experience.....	178
7.5 - Chapter Conclusion.....	181
8 - Conclusion.....	183
Bibliography.....	189

Acknowledgements

First and foremost I offer my sincere gratitude to Professor Michael Wheeler for providing regular and extensive feedback on my work and general support throughout nearly four years of PhD study. I could not have hoped for a better primary supervisor and feel enormously lucky to have had the benefit not only of his vast expertise but also his fantastically conscientious and supportive approach to supervision.

I am grateful to my secondary supervisor, Dr Simon Prosser, for his efforts in looking over draft chapters of this thesis and feedback which has provided a different and useful perspective. Thanks are also due to the other faculty members who have been kind enough to provide comments on my written work. These include Dr Zoe Drayson, who has provided helpful feedback both during and outside of my annual review, and Dr Adrian Haddock, who led two of my annual reviews and has been valuable in encouraging me to more clearly articulate my understanding of the sensorimotor theory.

I would like to thank my fellow students and all the staff members at the St Andrews and Stirling philosophy graduate programme, particularly those based at the philosophy department in Stirling, for providing a friendly close-knit environment and exposure to so many different ways of doing philosophy.

Last but not least I thank my parents, Erika and Julian Silverman. In particular, I am grateful to my father for setting off my interest in philosophy of mind and consciousness after introducing me to the work of Searle, and both for providing moral and often financial support during my almost never-ending years as a student.

1 - The Sensorimotor Theory of Perceptual Experience

The sensorimotor theory is an influential, non-mainstream account of perception and perceptual consciousness intended to offer an improvement on orthodox theories. It is often taken to be a variety of ‘enactivism’, and in common with enactivist cognitive science more generally, it de-emphasises the theoretical role played by internal representation and other purely neural processes, giving theoretical pride of place instead to interactive engagements between the brain, non-neural body and outside environment.

Although orthodox cognitive scientific accounts of vision vary in their details, they commonly descend from the hugely influential work of Marr, brought together in his posthumously published book *Vision* (1982). Marr’s contention was that vision is a form of information processing, beginning with the retinal image and resulting in a detailed, neurally-realised representation of the outside environment, the deployment of which is sufficient for visual awareness. Allied to this general approach is a naive and, as we will see, implausible conception of vision, in which this representation takes the form of an expansive and richly-detailed topographic model of the outside world, the activation of which is sufficient for a richly-detailed visual experience.

Although orthodox theories never explicitly endorse the naive conception, it is not disowned by Marr’s basic claim that vision is a process of constructing a detailed internal representation based on the more limited information available from the retinal image. This has prompted sensorimotor theorists to suggest that research on vision is in many cases tacitly motivated by the naive conception, causing it to generate faulty results. Part of the sensorimotor theory’s ambition, in offering an alternative to Marr’s basic thesis, is to provide a theoretical framework which avoids conceptual mistakes and enables more fruitful empirical research. To this end, the theory claims that perception is not a process of representation, but of skill-driven bodily exploration of the outside environment.

In addition to offering a distinctive account of the processing that underlies perceptual consciousness, the sensorimotor theory aims to explain perceptual consciousness itself, in particular its phenomenal character and the relation between phenomenal and physical properties. Many people continue to be moved by anti-physicalist intuitions about

‘qualia’, and the sensorimotor theory aims to respond to these by providing compelling and principled reasons for thinking that there is nothing more to being in a given phenomenal state than being engaged in a particular kind of embodied interaction with the environment. As before, the theory is supposed to be an advance on accounts of consciousness that appeal solely to facts about the brain, which sensorimotor theorists think lack the resources to defeat anti-physicalist intuitions.

Besides responding to this issue, the sensorimotor theory provides a number of distinctive insights into the logical and phenomenological character of perceptual experience. The theory, interestingly, integrates its contribution to various non-empirically oriented debates about perceptual experience with its account of the cognitive science of perception. In addition to taking a critical stance toward internal representation in cognitive science, the theory rejects representationalism in a number of other domains, including the attempt to explain phenomenal qualities by appeal to representational content, and representational realist accounts of the epistemology and metaphysics of perception. In place of representation, the theory appeals in each case to a bodily relation the perceiver stands in to the outside environment.

The theory was given its best known early statement in a paper by O’Regan and Noë, titled ‘A Sensorimotor Account of Vision and Visual Consciousness’ (2001a) - in fact, it is an account of perception more generally, although the focus is on vision. I will henceforth use the abbreviation O&N to refer to that paper or the paper’s authors in cases where it is that particular paper I am referring to. The other best known statement of the sensorimotor theory is Noë’s (2004) book *Action in Perception*, which is joined by more recent book-length accounts by O’Regan (2011) and Noë (2012), and a number of early papers on the theory, frequently authored in collaboration with Myin (e.g., Myin and O’Regan, 2002; O’Regan, Myin and Noë, 2004; 2005). Although O&N is the first official statement of the theory, it bears strong similarities to an earlier account by Hurley (1998). O&N were not informed by it when they first formulated the account, but Noë went on to collaborate with Hurley (Hurley and Noë, 2003; 2007), suggesting a case can be made for considering Hurley’s book the earliest work in the canon. Other important early work by the theory’s originators, although non-canonical, includes O’Regan (1992), Pessoa, Thompson and Noë (1998), and Thompson

and Noë (2004). The above accounts, and others by the theory's instigators, are supplemented by a large body of secondary work on the theory.

The theory has had an impact in a range of sometimes only loosely-related disciplines. In particular, it has inspired a good deal of increasingly prominent theoretical and practical work in cognitive science: widely-discussed glosses on the theory from the very recent past include, for example, a dynamical systems model by Burhmann, Di Paolo and Barandiaran (2013) and a predictive coding account by Seth (2014). The theory has attracted interest and no little controversy among philosophers of cognitive science, prompting prominent lines of criticism, for instance, by Clark (Clark and Toribio, 2001; Clark, 2006; 2008; 2009), Block (2001; 2005; O'Regan and Block, 2012), Dennett (2001), Prinz (2006), and many others. It has also attracted attention from philosophers of perception, a prominent example of which is a published symposium featuring responses to the sensorimotor theory by Campbell (2008), Martin (2008) and Kelly (2008).

1.1 - 4E Cognitive Science and the Varieties of Enactivism

The sensorimotor theory is situated amongst a complex web of theoretical ancestors and close cousins, and I will draw attention to some of these later in the thesis, in cases where it is helpful to do so. For now, it will be useful to do just a little basic scene setting. Although the theory has had an impact in non-empirically oriented debates about mind and perception, it is primarily a philosophically-motivated entry into the cognitive science of perception and perceptual consciousness. In particular, it enters the running alongside other 4E accounts, a broad class of approaches that depart from the view, mainstream in recent decades, that cognition is best explained by nothing more than the brain's construction and deployment of internal representations.

4E cognitive science claims that cognition is properly characterised using one or more E-words: embedded, embodied, extended and/or enactive. Embedded cognition is cognition that must be explained by appeal to non-trivial facts about the environment in which the cognitive tasks are carried out. Embodied cognition is cognition whose architecture

incorporates the non-neural body as well as the brain.¹ Extended cognition is cognition whose architecture also incorporates the external environment. Enactive accounts frequently reject internal representation altogether, and claim that cognition is ‘enacted’ through a process of interactive engagement between the brain, non-neural body and outside environment. It is worth noting that the degree to which the competing 4E approaches can be reconciled or, alternatively, constitute quite separate and incompatible traditions, is a matter of debate: see, e.g., Ward and Stapleton (2012) for an argument that they can be brought together, and Chemero (2009) and Wheeler (2014) for arguments to the opposite effect.

As I have mentioned, the sensorimotor theory is often regarded as a kind of enactivism. The best-established version of this approach is the ‘biological’ variant, notable statements of which include Varela, Thompson and Rosch (1991) and Thompson (2007). Enactivists in this tradition combine insights from phenomenology (in particular, Merleau-Ponty, 1945/2013) with work on the biological phenomenon of autopoiesis (e.g., Maturana and Varela, 1980). They claim that the interactive processes associated with cognition are continuous with those associated with life. Thompson observes that a cell, for example, produces its own material constituents through a metabolic interaction with the outside environment. Through this process of self-creation, the cell establishes the significance of certain environmental perturbations as either leaving the organism intact or resulting in its disintegration, and in this sense creates its own environmental niche. Living organisms arise out of self-organising biological processes like these in the way a tornado arises out of a dynamical process: they are, as Thompson puts it, invariant patterns that subsist in a changing physical substrate, implying they are not merely causally dependent on interactive processes, but constituted by them. Biological enactivists claim that self-organising biological processes of an appropriate kind are by nature minded, and that appealing to them helps close the explanatory gap between consciousness and the physical. To count as minded, an organism must be autonomous and answerable to norms. The appeal to biological self-organisation helps explain this. Organisms are autonomous by virtue of self-organising, and answerable to

¹ The label ‘embodied cognition’ is also sometimes adopted by quite different accounts, in which cognition is wholly brain-bound but involves neurally-realised representations of bodily states.

norms because their existence is predicated on responding to the environment in an appropriate way.

A recent entrant to the enactivist running is Hutto and Myin's (2012) 'radical enactivism'. In common with most statements of the sensorimotor theory, the authors exclude the biological apparatus distinctive of enactivism in its original form. Their distinctive contribution is to emphatically reject any theoretical role for representation or content, i.e. physical states bearing truth or correctness conditions about other physical states. In place of content, the authors claim that 'basic' cognition, meaning any cognition that is not shaped by natural language, consists of nothing more than interactive engagements between the brain, body and the external objects to which the cognition is intentionally-directed, which the authors propose can be explained by appeal to a nonrepresentational teleofunctionalist theoretical framework.

Noë gave the label 'enactive' to his (2004) version of the sensorimotor approach, not to declare a strict allegiance with biological enactivism, but to highlight its resemblance to it. Like biological enactivism, the sensorimotor theory downplays internal representation, and emphasises the role played by bodily interaction with the environment. In particular, it claims, like biological enactivism, that perception depends on the perceiver's mastering, in a bodily way, the character of its own bodily relation to the environment, in particular the way its orientation to objects in the environment would change given possible movements. Noë also expresses some sympathy toward the enactivist claim that consciousness requires life.² Thus the sensorimotor theory is typically regarded, with some justification, as a variety of enactivism. Hutto and Myin complain that the sensorimotor theory retains objectionable orthodox commitments on the ground that it retains a commitment to content and representation (although whether they are right about this is controversial, and Myin is himself, we should note, one of the original proponents of the sensorimotor theory). Despite criticising the sensorimotor theory for being insufficiently radical, the authors have, by including it in their taxonomy of enactive theories, helped crystallise the consensus that the sensorimotor theory deserves the 'enactivist' moniker.

² This claim does not, however, feature elsewhere in the sensorimotor theory, and is rejected by O'Regan, 2011.

1.2 - Structure of Thesis

In spite of its influential nature, the sensorimotor theory is formulated in differing and sometimes inconsistent ways throughout its canonical statements as well as the secondary literature, and in some cases its commitments have not been spelt out in sufficient detail at all. For instance, it is not clear whether the theory should, broadly speaking, be regarded as partly continuous with representationalist approaches in cognitive science, or allied with anti-representationalist embodied or enactive approaches such as those endorsed, for example, by Thompson (2008), Chemero (2009) and Hutto and Myin (2012). More narrowly, it is not clear what the sensorimotor skills the theory appeals to actually are, nor what kind of physical mechanism is implicated in them. Other examples of vaguely or inconsistently defined theses will become evident as the thesis progresses, and part of my aim is to identify aspects of the sensorimotor theory that are poorly defined and go some way toward clarifying them.

In chapter 2, I attempt a careful formulation of the negative claims that properly motivate the sensorimotor theory, an aspect of the account which is easily misunderstood and which I argue has been misunderstood, in one notable case, by Prinz (2006). Chapter 3 serves the dual purpose of laying out the sensorimotor theory's tenets in detail, and offering what I take to be the most promising and internally consistent ways of understanding them. Chapter 4 identifies some key limitations with the theory's notion of sensorimotor 'understanding', and offers a new proposal about what we should take sensorimotor understanding to be.

A later part of the thesis addresses some of the key criticisms that have been levelled against the theory, in particular by Clark (2008) and Block (O'Regan and Block, 2012). Chapter 5 addresses a worst case scenario in which it has been established that perception does in fact depend on nothing more than the activation of internal representations, which are not themselves intimately involved in the agent's sensorimotor engagements with the outside environment. I argue that in this event, there is still good reason to endorse the sensorimotor theory as an account of phenomenal qualities, and I offer some suggestions about how such an account could work. Chapter 6 offers a more robust defence of the sensorimotor theory in which I attempt to make it plausible that perceptual experience is constituted by the exercise of capacities for skillful bodily behaviour, and argue that representations should at most be

thought to play a heavily restricted theoretical role.

Chapter 7 moves beyond the core principles of the sensorimotor theory and begins to address new contributions the sensorimotor theory can make to long-discussed perceptual phenomena. Specifically, it addresses temporal extension. I argue that both the activity of perceiving and the contents of perceptual experience are temporally extended. Thus I suggest that object experience as well as event experience has a temporally-extended content, which temporally tracks the process of perceiving.

The aims of this thesis can be expressed as a unified goal. By clarifying some of the theory's founding principles, I aim to show how it can address some of the most prominent objections that have been levelled against it. By doing this, I hope to illustrate the novel contribution the sensorimotor approach can make to resolving some longstanding puzzles. Thus I hope to leave the sensorimotor theory in better shape than I found it.

2 - Making Conceptual Space

The sensorimotor approach makes a number of striking and relatively novel claims about perceptual consciousness and the physical processes that underlie it. Because of the more or less unorthodox nature of its positive claims, it is natural that the authors aim to motivate the approach by highlighting, with some emphasis, the faults and limitations of older approaches, and large early sections of O&N and Noë (2004) are devoted to this. The problem with this negative aspect of the sensorimotor theory is that many of the claims under target are not widely endorsed by contemporary theories, and in some cases merely reflect naive misconceptions or defunct historical ideas. In its standard presentation, the sensorimotor account is not pitched as a rejoinder just to mistakes of this sort, but also to the well-established orthodox view of vision in cognitive science, which holds that perception is realised by the construction and deployment, in the brain, of a detailed representation of the outside world.

The main part of this chapter will examine the manner in which the objections levelled in Noë and O'Regan's original statements of the sensorimotor approach cast doubt on the mature orthodox model. I will argue that they do not do enough by themselves to displace the orthodox view's core dogma that visual consciousness depends on internal representation, but they do, at least, indicate that orthodox theories must be careful to avoid certain mistakes. They may also intimate that the orthodox model is founded, implicitly, on a misplaced ambition, namely the desire to repair an obsolete 'inner movie' idea of vision that would better be replaced entirely. Because of the limited reach of the authors' negative claims, the sensorimotor account must be framed so that its appeal does not entirely depend on them.

2.1 - Kepler and Descartes: The Lessons from History

In a passage titled 'Fallacies Pictorial and Homuncular', Noë (2004) addresses historical accounts of vision, with a view to drawing morals for present day theory. Drawing on Lindberg's (1976) overview of pre-modern accounts, he calls attention, in particular, to Kepler's work on the optics of vision. Given what he knew about the eye's anatomy and the

principles that govern the refraction of light, Kepler was able to correctly identify that light meeting the lens is inverted (from top to bottom) and reversed (from left to right) before finally being projected onto the retina. This finding troubled Kepler. In addition to noting, quite appropriately, that the retinal image plays an important role in conscious vision, he thought that it did so by actually functioning as a picture. If the retinal image is inverted and reversed relative to the outside environment, it does not bear a pictorial resemblance to the environment as perceived. For this reason, Kepler ‘tortured’ himself attempting to identify a mechanism by which the light is re-inverted and re-reversed before reaching the retina. Finally abandoning this effort, he concluded that the disparity between visual experience and the projected image had to be corrected by a visual faculty occurring beyond the retina - involving, as he thought, spirits - and falling outside the scope of his own interest in vision, which was limited to optics.

Noë notes that Kepler’s misconceived search for a mechanism to re-invert and re-reverse the light was motivated by the more basic error, which Descartes later cautioned against, of supposing that the retinal image requires a further perceiver, a homunculus, to look at it. This is a faulty way of thinking, since it merely introduces the requirement for another homunculus to explain how the first one sees, and, following this pattern, an infinite regress of homunculi prevents explanatory closure. We should note that it is not obviously incoherent, in principle, to imagine that an image functions as a picture even in the absence of an observer, although claiming that a mere image constitutes a ‘picture’ requires some kind of explicit justification. Consider that a naturally occurring and function-less image such as a reflection of the sky in a lake does not merit being described as a picture in normal usage. Kepler, in any case, did think a homunculus was needed to view the picture, and for this reason was guilty of attempting to explain a psychological trait by introducing a homunculus that possesses the very same psychological trait - the most elementary version of what Kenny (1971) has called the ‘homunculus fallacy’.³ In spite of warning against the error, Descartes committed it himself by implying that the immaterial soul, serving as a homunculus, registers images presented to it at the pineal gland.

³ In section 2.1.3 I will address some different kinds of homuncularity, but this for now is the kind of homunculus under consideration.

2.1.1 - Pictures, Filling-In and the Bridge Locus

Noë (2004) claims that a confusion similar to the one that dogged Kepler persists into the present day. He gives the example of filling-in, a process by which the brain is alleged to compensate for information missing from the retinal image. Proponents of filling-in claim that the brain makes inferences about environmental features obscured, for example, by the retina's blind spot, and uses this inference-making mechanism to reconstruct the missing detail in an internal representation (e.g., Goodman, 1978).

On this topic, Noë, with reservations, endorses Dennett's (1991; 1992) position. Dennett is known for criticising the view that the brain's visual processing is geared toward the production of a Cartesian theatre, an end-product which occurs in a spatially circumscribed area of the brain and constitutes the necessary and sufficient substrate of visual consciousness. Dennett claims that this view, 'Cartesian materialism', is motivated by a tacit commitment to the existence of a homunculus for whose benefit the visual processing takes place. This is a problematic idea for the reason I have just noted, although Dennett is not concerned with homuncular regress so much as with implausible theoretical views that follow from positing a homunculus, the full details of which need not be rehearsed here. He observes that Cartesian materialists are forced to presume that the brain fills-in data missing at the blind spot because they imagine that the notional homunculus would otherwise notice that the information is missing. Because he rejects this position, Dennett claims instead that the brain simply ignores the lack of information at the blind spot. Noë (2004) enthusiastically endorses Dennett's rejection of Cartesian materialism, and in diluted form the scepticism about filling-in. He makes the important proviso that filling-in is only a misguided posit if the need for it is assumed *a priori* as a result of prior commitment to the existence of a homunculus. This does not rule out the possibility that the existence of filling-in can be established by appeal to relevant empirical work.

Visual stability is a further domain in which Noë suggests contemporary theories are sometimes guilty of the homunculus fallacy. It is sometimes considered noteworthy that despite continual saccadic eye movements and attendant changes to retinal input, the world is presented in visual experience as being stable. There is a temptation to assume that the brain

must deploy a special mechanism to compensate for this, as if there were a homunculus who would otherwise confuse the resultant changes to the positions of objects in the retinal image for changes in the spatial position of the objects themselves. Noë argues that a compensatory mechanism is only a conceptual necessity if you think that perception consists of a pictorial representation derived from the retinal image and offered up for the benefit of a homunculus. If you do not think this, you need not imagine that visual experience is in the first place prone, without a special compensatory mechanism, to mistaking eye movement-related changes in retinal input for changes in the outside world. Thus, an *a priori* commitment to the existence of the compensatory mechanisms only arises if you commit the homunculus fallacy. O&N and O'Regan (1992) make a similar point concerning saccadic suppression. Saccadic eye movements create a motion blur effect on the retinal image, and it is sometimes thought that the brain inhibits informational input from the retina during saccades in order to prevent perceivers from experiencing a blurred visual scene (Matin, 1974). O&N suggest that this is likely to be motivated by a commitment to a homunculus, since the smear effect need not be compensated for if the retinal image does not function as a picture.

These examples aside, Noë is imprecise about the intended scope of the historical lesson. He claims that: “No contemporary theorist believes that we see by seeing internal pictures” (2004, p. 46), but finishes the passage by stating: “The purpose of this section has been to reveal the extent to which our thinking about perception, like that of Leonardo and Kepler, is tied to a problematic conception of the need for pictures in the head” (p. 48). What Noë means when he says that theories are *tied to* this problematic conception is ambiguous, given that there are many ways a theory could be ‘tied’ to some conception which need not involve explicitly endorsing it.

Prinz (2006), responding specifically to *Action in Perception*, takes it that Noë’s intention is to draw a rough and ready parallel between Kepler and a contemporary version of the ‘pictorial’ view, in which neural states yield visual awareness with a particular content by virtue of their pictorial qualities. Prinz observes that: “We do not have color filled pixels in the head, much less an inner observer to discern all the parts. But this view is not seriously defended by anyone” (2006, p. 12). Appealing to Kosslyn (1994) as more representative example of a contemporary pictorial theory, he states that Kosslyn takes care not to require an

infinite regress of homunculi, and observes that pictorial views in general are not pre-committed to the existence of filling-in at the blind spot, since they could instead endorse Dennett's suggestion that the subject simply ignores the absences of information.

Prinz continues by citing a study by Ramachandran and Gregory (1991) in which subjects were presented with a television screen showing at first a 'snow' pattern and subsequently switched to show a solid colour instead. Subjects reported undergoing a briefly continued awareness of the snow pattern in an area of the visual field corresponding to the blind spot, even after the stimulus had been changed. The study suggests that information missing at the blind spot is not, in fact, simply ignored in visual processing, as if it were there would be no way to explain the continued percept of the snow in this region of the visual field. Prinz claims that the evidence that filling-in does take place counts as an embarrassment for the sensorimotor theory.

Noë is admittedly vague about the lesson he intends us to take from Kepler and Descartes, and a sceptic would be well-justified in observing that Noë does not comprehensively address the numerous resources that orthodox, including pictorial, accounts appear to have to avoid the homunculus fallacy. However, Prinz's rebuttal falls wide of the mark. Noë neither claims that pictorial accounts necessarily commit the homunculus fallacy and therefore fail, nor denies the existence of neural filling-in. A better account of his position can be found in an earlier paper by Pessoa, Thompson and Noë (1998). The paper rejects Cartesian materialism and also a doctrine the authors call 'analytic isomorphism', the *a priori* thesis that the substrate of perceptual consciousness must be a circumscribed area of the brain - a 'bridge locus' (Teller and Pugh, 1983) - which bears a functional (including topological) or topographical (i.e. map-like) isomorphic correspondence to the features represented by the subject's experience. Cartesian materialism and analytic isomorphism resemble one another in that they are both committed to the existence of a bridge locus, although isomorphic correspondence between the neural and the perceptual is in principle not essential to Cartesian materialism, and proponents of analytic isomorphism could deny that their position is motivated by any tacit or explicit commitment to the existence of a homunculus.

Pessoa et al. claim that if you endorse analytic isomorphism, you have a different reason to assume *a priori* the existence of filling-in. To maintain the isomorphism between the

bridge locus and the content of perceptual experience, there must be a neural state which is isomorphic to features that the experience represents as occurring in the location obscured by the blind spot, given that the subject does not experience the blind spot as an absence. The authors take it that once we reject Cartesian materialism and analytic isomorphism, there is no reason to aprioristically assume the existence of filling-in, although they actively argue for the claim that filling-in does take place, appealing to empirical data which I will not rehearse given that there is no disagreement on this point between the authors and Prinz.

A similar point holds with regard to pictures in the head. The pictorial view states that image-like physical states help generate visual awareness by serving as pictures, and this implies they are topographically or topologically isomorphic to the features represented in the experience, regardless of whether or not there is supposed to be a homunculus to look at them. Thus, pictorial theories could fall under the doctrine of analytic isomorphism, although they need not. Pessoa et al. are explicit, however, that they are open, subject to empirical evidence, to the claim that neural states help enable visual consciousness by virtue of bearing a pictorial resemblance to the features represented in visual consciousness. The authors reject the aprioristic commitment to isomorphism, along with the suggestion that where pictorial representations do play a role, this takes place in a spatially localised neural region which serves as a bridge locus. Noë (2004) says nothing to contradict any of the claims I have just mentioned from the (1998) co-authored paper.

Notice that for each of the phenomena Noë (2004) and Pessoa et al. (1998) target - pictorial representations, filling-in, compensatory mechanisms for visual stability, and neural-perceptual isomorphisms - the claim is not that these phenomena do not exist. The authors are careful to claim that the phenomena may exist, and in some cases probably do. A better moral is that confusion arises when we make theoretical commitments to Cartesian materialism or analytic isomorphism necessitating the existence of something that should be an empirical posit, a diagnosis that at least holds true when we think of Kepler's commitment to pictorial representations and a homunculus. In one sense, then, the negative claims motivating the sensorimotor theory are more modest than Prinz suggests. Noë should not be taken to be attempting to actively disprove pictorial theories or other representational theories that, like Kosslyn's, take care not to commit the homunculus fallacy.

Prinz's other objection is that Noë is targeting a straw man, given that no one seriously defends the claim that we have colour-filled pixels in the head or a perceiver to discern them. Strictly, this may be true. Pessoa et al. observe, however, that in the commentaries to their *BBS* paper, Lehar (1998) explicitly claims that spatial perception logically requires the existence of a three-dimensional neural model. In his own widely-cited *BBS* paper, Lehar (2003) claims that topographical or topological isomorphism is not a logical requirement for perception, but that functional isomorphism is. Using the (2003) paper as a guide, Lehar's requirement can be reconstructed in the following way. Imagine that the brain contains a 3D model in the shape of a cube and that this correlates with the conscious percept of an identically-shaped cube. The model is thus topographically isomorphic to the perceived cube. Now suppose that instead of being perfectly cube-shaped, the model has the shape of a stretched-out cube, with some planes that are oblong rather than square. The model continues to be topologically isomorphic to the perceived cube, because the model's planes continue to stand in appropriate spatial relations to one another. Informally, this is because a cube could be distorted into a stretched-out cube shape without ripping or tearing.

Lehar claims that topological or topographical isomorphisms are likely to play a role in perception, but are not required as a matter of logical necessity. Now, imagine that instead of merely being stretched-out, the cube-shaped neural model is scrambled at random, such that it no longer looks anything like a cube, even a distorted one. This means it is no longer topologically or topographically isomorphic to the perceived cube. Think of the scrambling, however, as a warping of the geometry of the space the shape inhabits, rather than merely a warping of the shape itself. Suppose that the perceived cube is transformed in some way, for example rotated. If the model is transformed in the way that a topographical model would have been, allowing for the distorted geometry of the space it inhabits, a functional isomorphism obtains between the model and the perceived cube.⁴

Lehar insists that functional isomorphism, at least, is logically necessary for perception. This is a strikingly robust claim, and almost sufficient for the theory to constitute an instance of analytic isomorphism, although the appeal to functional isomorphism avoids the logical

⁴ Pessoa et al. (1998) use the term 'topological isomorphism' to refer to topological *or* functional isomorphism in Lehar's usage. I am following Lehar's usage.

requirement for a localised bridge locus. Lehar suggests, however, that there is reason to think, prior to empirical confirmation, that topographical and topological isomorphisms are particularly likely to obtain, since they are more efficient than other functional isomorphisms to implement. This is a demanding theoretical commitment, and it suggests that Lehar's view is biased toward the existence of a bridge locus, even if it does not require it, since these allegedly more efficient forms of isomorphism cannot be distributed across the brain in the way functional isomorphisms could be. This shows that analytic isomorphism is not far removed from explicit claims found in recent and influential work.

Lehar's view is not a mainstream position, however: A moral more widely applicable to representationalist cognitive science is that we should be deliberate about rejecting any aprioristic commitment to posits that are only needed if one does endorse analytic isomorphism or Cartesian materialism. Notice that doing so opens up conceptual space for a range of posits that, unlike pictorial representations and filling-in, are actively incompatible with those commitments; and some of the alternative views thereby made available may prove persuasive on positive rather than negative grounds. After rejecting Cartesian materialism, Dennett claims that consciousness depends instead on representational processes that are widely distributed across the brain. The sensorimotor theory does something similar, but makes a bolder positive proposal by claiming that cognitive processes sometimes depend in a non-trivial and frequently unacknowledged way on the subject's embeddedness within the outside environment and, on some formulations, that the substrate of perceptual experience actually includes activity in the extra-neural body and outside environment.

I will elaborate on the sensorimotor theory's positive claims in chapter 3, but will here briefly highlight the sensorimotor theory's alternative proposals with regard to filling-in. Pessoa et al. observe that rejecting analytic isomorphism and Cartesian materialism allows for a pluralistic explanation of perceptual completion, the visual phenomenology sometimes accounted for by filling-in. They propose that in some cases this involves pictorial representation, in other cases symbolic representation, and in others requires the absence of information to be ignored in the way Dennett suggests. Their account also holds open the possibility that perceptual completion might in some instances be explained by appeal to O'Regan's (1992) idea of the world as an 'outside memory'. O'Regan's idea is that perceptual

experience depends on a temporally-extended, embodied process in which perceivers gain access to missing detail by attending to different parts of the visual field, making use, in particular, of eye movements to access missing detail as needed directly from the outside environment without needing to 're'-present that information internally (see section 3.1).

2.1.2 - Neural Representation

I now want to offer another lesson we can learn from the historical discussion featured in Noë (2004), in a similar spirit to arguments made by Noë (2004) and Pessoa et al. (1998) but more unorthodox in its implications. Just as it would be wrong to assume, without good evidence, that filling-in takes place, it would also be a prejudice to commit oneself in advance to the claim that the neural processes involved in perception - including any inner processes that compensate for information missing from the retina - must involve representation. Some theorists (e.g., Wheeler, 2005) claim that internal representation requires homuncularity, although not necessarily of a problematic kind. I endorse this claim in section 2.1.3. The important point for now, however, is that to assume out of prejudice that the neural states involved in perception involve representation is to betray a commitment to the kind of homunculus which Descartes warned against. Descartes observed that retinal images play a causal role in perception, but that you are likely to be in the grip of the homunculus fallacy if you conceive of them as 'pictures'. By the same token, the states and processes involved in perception, including those involved in filling-in, might play a causal role without functioning as representations. It would be a mistake to pre-theoretically commit oneself to the existence of representation when the neural activity that helps enable perceptual experience might be better conceptualised as a nonrepresentational causal process.

Deciding whether the neural states involved in perception, including filling-in, are representations is in part an empirical enterprise. Indeed, it is possible that appeals to internal representation could be justified on entirely empirical grounds. Suppose that there were a unified structure in the brain that looks like an expansive and finely-detailed three-dimensional model of the visual scene, and we could use this to accurately predict the character of the subject's reported perceptual state. Few theorists realistically entertain this,

and even Lehar (2003) falls short of insisting on it, but it remains a conceptual possibility. If this were the case, it would be reasonable to suppose that the neural structure is a 'model' or 'representation', even in the absence of a philosophically robust explanation of how the structure comes not just to resemble the world but, strictly speaking, to 'represent' it.

A similar possibility holds true even if the purported neural representations are not supposed to involve neural-perceptual isomorphisms. For instance, it might turn out that we can accurately predict a subject's perceptual state on the basis of the subject's present neural state in combination with a theory of vision that makes reference to representation purely because appealing to representation is the most useful way to explain the neural activity. The idea that a neural state can be legitimately described as a representation purely because it is expedient to do so is widely endorsed, even by non-representational theorists like Chemero (2009), although it is rejected by Hutto and Myin (2012) - more on which in section 6.4.

If we agree that there is filling-in, this does not licence the claim that there is necessarily representation. There may be processes that complete visual information missing from the retina and in so doing affect the character of visual consciousness. However, should we abandon the view that there is a detailed inner world model (a kind of representation) to begin with, we abandon another important motivation for insisting that these processes necessarily constitute representation, even of a non-pictorial kind. The suggestion that there are processes that resemble filling-in without representation should be no more controversial than Descartes' assertion that the retinal image plays an important causal role without functioning as a picture. If we construe the import of the Kepler discussion for the sensorimotor theory this way, we have another reason for thinking that Prinz's (2006) contention that empirical evidence for filling-in should embarrass the sensorimotor account is wide of its mark.

It is an empirical bet, however, that neural states best understood as representations do exist. If neural representations do exist, it is also an empirical bet that appealing to their content or functional role is sufficient by itself to explain the target phenomena. For example, we might find that some neural process plays a functional role equivalent to filling-in, but also find that the process, taken together with other neural processes explicable in terms of representation, are insufficient to show which features the perceiver is intentionally-directed

toward at the agent level, or to account for the attendant phenomenal character (if the intentional directedness and phenomenal character of an experience could come apart, a possibility that may not be coherent). Pessoa et al. (1998) and Noë and Thompson (2004) deny that subpersonal content can be used to predict personal level perceptual states in this way, although they do not deny that subpersonal representation exists.⁵ If it proved impossible, however, to explain or predict the intentional or phenomenal character of the subject's experience without pointing to non-neural features, including, for example, variations in the agent's behavioural capacities, this would deprive representationalists of a significant reason for thinking that perception is realised by internal representations at all.

The point of this passage has not been to claim that perception does not involve representation, but, as before, to make conceptual space for the possibility that internal representation might play a reduced or eliminated explanatory role in a final account of perception and phenomenal consciousness. Many who are open to representationalist

⁵ Personal-level mental states are those it is appropriate to attribute to a whole person, and are generally thought to necessarily be conscious or at least consciously accessible. Subpersonal states are those which it is appropriate to attribute to a person's parts. For example, a neural representation may bear content about some state of affairs, even though the conscious subject herself does not entertain mental states with that content. Some states may be subpersonal and personal, just in case they can be intelligibly attributed to a subpersonal subsystem but are also consciously accessible. Noë (2004) expresses doubt about the value of the distinction on the ground that certain activities, such as eye movements, sometimes come under rational control by the person, and sometimes do not, and that both cases are likely to be similarly important to perception. What he has in mind is McDowell's (1994b) use of the distinction, which Drayson (2014) implies is an attempt by the Pittsburgh school of philosophers to hijack the distinction for their own idiosyncratic purposes. Personal level states on this perspective are those occurring in the space of *reasons*, meaning they are properly considered to be mental states, while subpersonal states are those occurring in the space of *causes*, the latter category denoting states that are not genuinely mental states at all but non-minded physical states. Note that on the non-Pittsburgh understanding, the fact that subpersonal states are sometimes assimilated to the personal level is to be expected, meaning that Noë's point does not necessarily compromise the distinction in its original meaning. In any case, Noë's sole-author account is marked by notable shift from talking about brains and artifacts, as O&N did, to talking about whole persons. Moreover, he endorses a distinction between the 'constitutive' and the merely 'enabling', which in McDowell's paper tracks the distinction between the personal and subpersonal, and which may have a significance analogous to the personal/subpersonal distinction in Noë's account. I explain this latter distinction and examine its significance later in the thesis.

accounts of cognition endorse the suggestion that the legitimacy of internal representation as a theoretical posit cannot just be assumed, but must be justified relative to the particular theory it is deployed in (see Ramsey, 2007, for a detailed defence of this point). However, others who are critical of non-representational accounts are prone to claiming that cognition, including perception, *just is* a representation-involving process.⁶

2.1.3 - Varieties of Homuncularity

I have, up until now, been discussing the most elementary version of the homunculus fallacy, identified by Descartes when he rejects the suggestion that retinal images function as pictures, warning: “We must not think that it is by means of this resemblance that the picture makes us aware of the objects - as though we had another pair of eyes to see it, inside the brain” (1637/1985a, p. 167). Ironically, contemporary theories that make this kind of mistake are described by Dennett (1991), pejoratively, as examples of Cartesian materialism, on the ground that Descartes went on to make the same mistake by claiming that the immaterial soul is presented with images at the pineal gland.

This elementary version of the fallacy is considered by some (although not Dennett) to be just one instance of a wider fallacy, allegedly more pervasive than mere Cartesian materialism in present day cognitive science. The error, sometimes known as the mereological fallacy, can be defined, broadly, as the ascription to parts of things predicates that should properly only be ascribed to wholes. Narrowly, it is the ascription of psychological predicates, for example ‘seeing’ or ‘believing’, to parts of persons - in particular, brains - where they should only be ascribed to persons, i.e. the complete living organism (Kenny, 1971; Bennett and Hacker, 2003). In the basic version of the fallacy, applying a psychological predicate to the brain is uncontroversially a mistake. If we try to explain the faculty of vision by positing the existence of a homunculus (a part of the person) that possesses an identical faculty, we find an infinite regress ensues of the kind discussed earlier. It is telling that the fault, in this instance,

⁶ For proof that many theorists take cognition to involve representation *by definition*, see Hutto, Kirchhoff and Myin (2014), who offer textual evidence from Khalidi (2007) and O'Brien and Opie (2009) to this effect.

lies in the circularity of the proposed explanation.

A.O. Rorty (1971), replying to Kenny (1971), notes that circularity of this sort only occurs if you try to explain a person's psychological state by ascribing to their brain or parts thereof *exactly the same* psychological state. The threat of circularity does not indicate the existence of a more general problem with ascribing person-like psychological predicates such as 'believing' to brains, since circularity is avoided so long as the explanandum and explanans feature different psychological states. Dennett (1978) is mindful that a certain kind of circularity may persist, however, even if the intentional predicates attributed to the part and the whole are different, or attributed to different degrees. To account for individual faculties like thought or vision, we need to explain, more generally, how intentionally-directed psychological states can be realised by a physical system at all. To account for a psychological attribute possessed by the person by reference to psychological attributes possessed by the brain is to presuppose that something bears intentionality, even if the states or processes featured at each end of the explanation are different. Dennett therefore proposes, like Rorty, that a person with complex psychological attributes might subdivide or 'decompose' into parts - i.e. homunculi or 'subpersons' - possessing simpler psychological attributes. Those homunculi might, in turn, decompose into still simpler homunculi. To show how psychological attributes might be realised in the first place, Dennett introduces the crucial suggestion that the bottom layer of homunculi are so simple, they decompose or 'discharge' into physical processes that can easily be described using non-intentionalistic vocabulary.

For some, however, circularity of the sort I have been discussing is merely a symptom of a more basic error, a variety of category mistake. Bennett and Hacker (2003) and Kenny (1971) both quote, approvingly, Wittgenstein's assertion that: "Only of a human and what resembles (behaves like) a living human being can one say: it has sensations; it sees; is blind; hears; is deaf; is conscious or unconscious" (1953/2009, para 281). The moral, for Kenny, is that since it is not *obvious* that brains resemble humans in the relevant way, we should proceed with caution, rather than unthinkingly ascribing the properties actually possessed by the whole person to parts thereof. Bennett and Hacker (2003) draw a more demanding conclusion. They argue that it is evident from everyday linguistic convention that psychological or otherwise intentionally-loaded vocabulary such as 'representing', 'mapping',

‘understanding’ and ‘believing’ can only be meaningfully used in relation to persons rather than brains. On this basis, and irrespective of any particular empirical facts about brains, Bennett and Hacker claim that neuroscientists, psychologists and cognitive scientists are guilty of conceptual incoherence when they apply these predicates to the operation of brains.

While homuncular decomposition offers, in principle, a means to avoid circularity, it does not, in itself, address the more fundamental charge that brains and parts of brains simply do not, in the relevant way, resemble or behave like humans. Dennett (2007), replying to Bennett and Hacker, remarks that we cannot use ordinary linguistic practice alone as a guide to the applicability of intentionalistic predicates to subpersons, since the claim that we can use these predicates turns on the empirical thesis that subpersons do, indeed, behave like human beings. Dennett is surely right that the admissibility of homuncular decomposition depends not just on aprioristic conceptual analysis but also on empirical facts about brains. There is a pressing uncertainty, however, about the sort of properties or behaviours that would need to be identified through empirical enquiry to show that an object is sufficiently person-like to licence the ascription of psychological predicates.

Dennett (1987) offers a specific proposal that could answer this. According to the ‘intentional stance’, we can legitimately attribute psychological predicates like beliefs and desires to an object, including a person, creature or artifact, just in case ascribing them is a useful way to explain and predict the object’s behaviour. The account is explicitly intended to be liberal enough to allow us to ascribe intentionality not just to people but non-living artifacts like thermostats and chess-playing computers.⁷ Dennett asserts not only that the intentional stance is scientifically respectable, but also that it reflects everyday folk

⁷ Dennett (1987) does not actually claim that we should apply the intentional stance to brains. He suggests that at this level we should apply the design stance, meaning we conceive of them as being designed to achieve particular functions. However, the homuncular decomposition strategy requires us to view brains and parts of brains as person-like, meaning the design stance does not appear to do the job. The intentional stance actually undercuts one of the roles homuncular decomposition might play, namely to resolve the fundamental question - a sort of hard problem - about how mere physical states could be intentionally-directed. The intentional stance solves this problem without requiring homuncular decomposition. However, endorsers of the intentional stance may still usefully make use of the homuncular decomposition strategy to show how intentionalistic descriptions can finally be cashed out using non-intentional, physicalistic vocabulary.

psychological practice, noting that in everyday conversation we frequently use belief and desire talk, for example, to make sense of behaviour by machines. If this is correct, then it is not implausible that brains, and parts of brains, can be ascribed intentionality in this way, too. In this event, the mereological fallacy is no threat to homuncular decomposition, either by the standard of everyday linguistic practice *or* the standard demanded by our theoretical account of intentionality.

Bennett and Hacker (2003) are on highly contentious ground when they suggest that everyday usage gives us a determinate guide to the correct ascription of intentional predicates, since one might think that we do not necessarily know precisely which features make our everyday ascriptions of psychological predicates true or false. It might be that in folk psychological discourse, we are employing the intentional stance. In this event, Bennett and Hacker are hoist with their own petard, for the reasons I have just suggested. Alternatively, it may be that everyday use of intentionalistic language picks out, unbeknown to us, properties unique to people or other complete living organisms, meaning the ascription to intentionalistic predicates to brains or artifacts are homonymic, metaphorical⁸ or just false. Bennett and Hacker's argument is enough to show that endorsers of homuncular decomposition have a positive need to endorse an account of intentionality that makes the approach tenable.

Depending on the approach to intentionality you endorse, the empirical criteria required to establish the existence of 'homunculi' may be more or less demanding. Adopting the intentional stance, homuncular decomposition seems relatively unproblematic. Alternatively, if you endorsed, for example, the autopoietic enactivist thesis that only a certain kind of autonomous biological agent can manifest intentionality (e.g., Thompson, 2007), homuncularity would appear much less tenable. Similarly, those representationalists that suppose intentionality is best explained by robustly naturalised subpersonal content (e.g., Millikan, 1984; see section 6.4) must maintain that the homunculi endorsed as 'top down'

⁸ McDowell (1994b), for instance, argues that content-talk at the subpersonal level is metaphorical, but nonetheless legitimate. In chapter 7, I argue that metaphorical representations and content is legitimate in principle, but may not be useful or appropriate in practice.

theoretical posits by Dennett's approach find counterparts in evolved modular architecture concretely implemented by the brain, a point made by Wheeler (2005). In either of the latter two cases, the burden of empirical evidence required to justify ascribing psychological predicates to parts of the brain will be strong.

I will not attempt to rule, now, on the correct account of intentionality. I will conclude this section by noting, however, that while Bennett and Hacker's automatic dismissal of intentionalistic vocabulary at the subpersonal level is misplaced, endorsers of this vocabulary must offer an account of intentionality that is amenable to the existence of homuncular decomposition, naturalistically plausible, and, we might think, motivated by considerations other than the desire to offer a *post hoc* justification for their use of intentionalistic language. This is a heavy burden, and it may prove that a more promising approach is to reformulate the vocabulary in question instead.

2.2 - Conservative and Radical Strands

As we have seen, the negative claims that properly motivate the sensorimotor theory are geared toward unseating a number of prejudices rather than actively ruling out the suggestion that perception depends on nothing more than the construction and deployment of internal representations, even representations that are richly detailed, pictorial or realised by a circumscribed area of the brain. However, the arguments are sufficient to create a legitimate suspicion that conceptions of vision endorsing these representationalist posits are motivated, at least in part, by a residual commitment to faulty historical assumptions, even if they have resources to avoid the homunculus fallacy.

The fact that the negative claims are relatively modest means there is leeway in how radical we take the theory's positive account to be. I will continue by outlining two competing versions of the sensorimotor approach, one *conservative* and one *radical*. Although they are potentially in significant disagreement with one another, the competing accounts reflect, I will argue, an uncertain attitude to representation present in the joint and sole-authored work of both Noë and O'Regan. The conservative account embraces a moderate construal of the authors' negative claims, and consequently endorses the orthodox thesis that perceptual

consciousness is realised solely by the brain (the claim known as *vehicle internalism*) through the construction and deployment of internal representations (*representationalism*); by virtue of this concession, the account can be thought of, essentially, as an attempt to progress the orthodox model in a way that inoculates it against Cartesian materialism, rather than an attempt to radically break from orthodoxy. This conflicts with the sensorimotor theory's frequently radical tenor. However, the conservative approach continues to make a number of interesting claims, including the view that perceiving depends on active, goal-directed interrogation of the outside environment.

Many of the theory's claims place a distinctive emphasis on embodiment: Although the conservative sensorimotor approach denies that perceptual consciousness supervenes on extra-neural activity, it recognises that bodily exploration plays a non-trivial and generally unacknowledged causal role in the construction of the neural states that do form the supervenience base. On this basis, it asserts that bodily exploration must be mentioned in any adequate account of vision. Further, the conservative approach endorses the particularly distinctive thesis that phenomenal qualities, rendered mysterious on rival accounts, are best explained by appeal to laws of sensorimotor contingency, as I will explain. This, again, means our account of perception must refer to characteristics of the perceiver's body, specifically its morphology and the nature of the body-environment relation that attends. By virtue of the body and environment's essential explanatory role, the sensorimotor approach could, on this construction, be described as a form of 'embedded' and conservatively 'embodied' cognition. Given the limited scope of the theory's negative claims, the relatively modest ambition of the conservative approach is an advantage, since it only needs to upset that claim that perceptual awareness depends on richly detailed, pictorial representations, and not the claim that perceptual awareness supervenes on internal representations in general.

An alternative construction situates the sensorimotor theory near the radical end of a continuum of 'embodied' and 'enactive' approaches to cognition. Accounts in this region frequently deny that cognitive processing is realised solely by neural states. To this effect, they sometimes suggest that the representational vehicles which realise cognition frequently incorporate bodily or environmental activity (Clark and Chalmers, 1998; Rowlands, 2006). Alternatively, they deny that there is representational content or anything else that would

motivate a principled insistence that the boundaries between brain and body, or body and world, are of theoretical relevance in determining the bounds of cognitive processing (Hutto and Myin, 2012). Accounts at the radical end of the spectrum, especially those that adopt the ‘enactivist’ label, are also known for suggesting that there is no internal representation whatsoever (e.g., Varela, Thompson and Rosch, 1991; Chemero, 2009; Hutto and Myin, 2012).

2.2.1 - The Tenets of Sensorimotor Theory: A Brief Outline

In this section I will proceed by outlining the main tenets of sensorimotor theory, categorising them into claims endorsed by both conservative and radical variants of the approach, and claims endorsed solely by one variant or another. I do not mean to claim that any version of the sensorimotor account must endorse every one of the claims falling under either the conservative or the radical strand. However, identifying these distinct variants offers a useful way to examine differences in the explicit claims made by the canonical accounts of the sensorimotor approach, different ways of reading those accounts, and disagreements in other literature on sensorimotor theory, a number of which will be investigated as the thesis progresses. The first four claims are endorsed by both variants of the theory and the claims endorsed only by one or the other follow.

(1) *The World as an Outside Memory*. This claim draws on work by O’Regan (1992) and ‘active vision’ approaches in computer vision (e.g., Ballard, Hayhoe, Pook and Rao, 1997). Endorsing Brooks’s (1991) suggestion, in robotics, that the world serves as its ‘own best model’, the outside memory thesis (O&N) denies that human visual systems make use of a detailed inner world model, claiming instead that our visual systems access detail directly from the environment as required in the course of bodily exploration. This claim is given weight by empirical work on change blindness and inattention blindness, where it has been found that human subjects, with their attention drawn elsewhere, frequently fail apparently to notice occurrences that take place seemingly in full view. The outside memory thesis entails that visual processing is, at least, causally scaffolded by bodily exploration in an interesting and frequently unacknowledged way. It does not necessarily imply the metaphysical thesis that

perceptual consciousness actually supervenes on extra-neural activity, although the outside memory thesis may be a key aspect of an account that does make this claim.

(2) *The Sensorimotor Contingency Account of Phenomenal Qualities.* Phenomenal qualities (or ‘qualia’), for example the feel of a sponge, the look of red or the sound of a bell (O’Regan, 2011), are not irreducible. They can be accounted for, exhaustively, by reference to characteristics of the conscious perceiver’s bodily relation with the outside environment - specifically, sensorimotor contingencies (SMCs). In O&N, SMCs characterise the law-like patterns of dependence that hold between motor output signals and informational input from the sense organs. That paper does not specify explicitly, but they might be counterfactual dependences, expressible by statements of the form ‘if I moved like this, sense input *would* be like this’. More restrictively, they might be patterns evident in the bodily interactions that the subject is right now engaged in, for example the way a sponge they are presently squeezing between their fingers yields easily to touch. Burhmann, Di Paolo and Barandiaran (2013) note that the latter kind of SMCs depend not just on peripheral sense input and motor output states, but on the internal states of the agent, since they depend on how the subject’s motor responses vary in line with different sense inputs, and this depends on the brain. The account of colour offered by many statements of the sensorimotor theory (O’Regan, 2011; Degenaar and Myin, 2013) also appeals to patterns of counterfactual dependence outside the skin, since an object’s experienced colour depends on the ways sense input changes with movement and in line with changes to ambient lighting conditions. Finally, in Noë (2004), SMCs also describe perspectival properties, the ways objects consciously appear when considered from a subject’s own perspective (for example, the elliptical appearance of a plate when viewed at a non-perpendicular angle) and the ways these change in line with movement.

(3) *The Sensorimotor Mastery Account of Perception.* Perceiving takes place when the subject possesses or exercises ‘knowledge’, ‘mastery’ or ‘attunement’ to SMCs. Sensorimotor ‘mastery’, as I use it throughout, is a generic term intended to incorporate a number of loosely related notions found throughout the literature on sensorimotor theory, and intended to be liberal enough to incorporate a ‘radical enactivist’ version of sensorimotor theory. The proponents of radical enactivism, Hutto and Myin (2012), dislike the terms sensorimotor ‘understanding’ and ‘knowledge’ on the ground that these notions appear to

require representational content. Myin and Degenaar (2014) endorse, by way of developing a positive radical enactive gloss, a notion of ‘attunement’ instead, which I include under the umbrella of sensorimotor mastery (note, however, that their concept of sensorimotor ‘attunement’ is different to the one I develop in chapter 3). O&N and O’Regan (2011) refer to a variety of mastery which appears, like attunement, to denote a capacity for skillful bodily behaviour that does not necessarily call upon internal representation. The same accounts suggest, however, that *conscious* perception requires sensorimotor ‘knowledge’, which might be taken to denote a distinct kind of faculty.

(4) *Perceptual Presence as Access*. This idea, endorsed by Noë, is a natural counterpart to the outside memory claim. It suggests that the environment, in all its detail and expansiveness, has a felt phenomenological ‘presence’ to the conscious perceiver which outstrips the information processed, at any one time, by the perceiver’s visual system. This is explained by the perceiver’s possession of sensorimotor mastery, which in this case amounts to knowledge of (or perhaps attunement to) the movements s/he must make to find out what is present and, perhaps, an implicit expectation about what the sensory results of those movements will be. The ‘presence-as-access’ (or ‘virtual presence’) the environment enjoys in the perceiver’s conscious awareness is analogous to the presence to your computer of the web edition of the New York Times. When you read an online newspaper, you generally download just one article at a time. However, since you can access any of the day’s articles the moment you require them, they are, practically speaking, just as ‘present’ as they would be if you downloaded the day’s edition all at once (Noë, 2004).

The following claims are endorsed only by the radical sensorimotor approach:

(5) *Virtuality All-The-Way-In*. The presence-as-access thesis, above, could be straightforwardly endorsed by most sensorimotor theorists as a natural phenomenological counterpart to the outside memory thesis. Noë makes a more challenging proposal, however, when he says that in a disanalogy with the computer case, perceptual presence in humans is virtual “*all the way in*” (2004, p. 134). This means that there is no distinction to be found in the phenomenal character of perceptual consciousness between detail that is now present

‘simpliciter’ and detail that is present by virtue of being accessible in the manner described by the presence-as-access and outside memory theses. Noë claims that this suggestion finds confirmation in the introspectable phenomenal character of one’s own visual consciousness, alleging that you can never take in, all at once, any complete or atomic visual feature, for example, “*that* shade of red” (p. 134). If the virtuality *all-the-way-in* thesis is considered primarily as an extension of the outside memory thesis, it might seem to amount to a denial that perceptual consciousness exists or has any phenomenal character. As I will later argue, this impression can be avoided if the claim is considered, instead, primarily as a counterpart to the next thesis.

(6) *Perception as Dynamical Entanglement*. This claim is endorsed by Noë (2004) and reflects, in particular, important work by Hurley (1998). It involves rejecting the suggestion, endorsed by many classical accounts of cognition (see, e.g., Fodor, 1983), that there are functionally distinct and informationally encapsulated ‘input’ and ‘output’ systems where the activity of input systems has no tight causal link with the activity of output systems, or vice versa. Endorsers of the dynamical entanglement thesis claim, instead, that there is an intimate causal link, running in both directions, between neural activity in the sense input and motor output areas of the brain. Further, the thesis claims that there is no theoretically relevant distinction between activity in those areas and outside, in the extra-neural body and environment. As a result, perception must be explained by reference to patterns of dynamic activity occurring across the entire loop of interaction occurring between brain, extra-neural body and world.

(7) *The Extensive Conscious Mind*. The dynamical entanglement thesis is one of the key claims motivating the closely allied thesis, endorsed by Noë (2004) and Hurley (1998), that the supervenience base of perceptual consciousness literally extends beyond the boundaries of the brain, incorporating aspects of the extra-neural body and outside environment. The argument, briefly put, is that since perceptual consciousness depends on a dynamical system incorporating these extra-neural features, it may not be possible or appropriate to distinguish between the contribution made by these features and the brain when determining the substrate. This thesis is sometimes described as ‘vehicle externalism’ about consciousness or the ‘extended conscious mind’, although for reasons that will be

explained in section 3.6, I settle on the more idiosyncratic term ‘extensive conscious mind’.

(8) *Anti-Representationalism*. Representation is a key concept in traditionally-styled cognitive science and features prominently in the orthodox approach to perception following Marr (1982). The sensorimotor theory is unequivocal in rejecting certain misplaced roles for internal representation. These include, in particular, the suggestion that a human perceiver possesses an inner model of the outside world sufficient to generate conscious awareness comprising, all at once, a large, richly detailed or picture-like visual field.

It is unclear whether the sensorimotor theory’s originators intend us to think that sensorimotor knowledge should be thought of as involving representational content. Noë (2004) explicitly declines to take a stance on whether or not cognition in general, including perception, involves or could involve neurally-realised representations. O&N and O’Regan (2011) say they are happy with talk of ‘representation’ in neuroscience when it is used, minimalistically, to describe nothing more than patterns of covariance between physical states. However, they stop short of insisting that there are or could be any representations carrying *content*, which must be evaluable by truth or success conditions, rather than just *information*, which need not, and they stop short of insisting that sensorimotor knowledge *necessarily* involves representation in any sense.

O’Regan (2011), after defending use of the term ‘representation’ in neuroscience, strikes an ambivalent tone by declaring that representation-talk is best avoided. The strongest sign that the sensorimotor approach may necessarily entail the presence of representational content occurs when O’Regan claims that perceptual consciousness depends on sensorimotor engagements plus higher-order thoughts about those sensorimotor engagements. O’Regan states, uncertainly, that the approach might have something in common with classic Higher Order Thought theories of consciousness which, we can note, do typically take it that consciousness depends on higher-order representations of lower-order phenomena (Carruthers, 2011). The sensorimotor theory’s uncertain attitude to internal representation has allowed some commentators to argue, approvingly or disapprovingly, that internal representation plays an essential role in sensorimotor theory as originally presented (Roberts, 2010; Hutto and Myin, 2012). I will later argue that the sensorimotor account, even in the form offered by its progenitors, could in principle dispense with any role in perception for

internal representation.

The conservative variant of the sensorimotor approach consists, largely, in adopting the core claims while rejecting the radical ones. It may also endorse, positively, the following tenets. I will list them quite briefly as they comprise, in essence, the converse of the radical claims:

(9) *Representationalism*. Exercise of sensorimotor knowledge is realised subpersonally by the deployment of internal representations.

(10) *Internalism*. Although perceptual processing is causally scaffolded by bodily exploration, perceptual consciousness is realised solely by the brain.

3 - The Tenets of Sensorimotor Theory

In chapter 2, I identified a number of theoretical posits toward which the sensorimotor theory takes a critical stance, including the bridge locus, pictorial representation, other forms of representation and, more narrowly, phenomena such as filling-in and saccadic suppression. Recall that the sensorimotor theory's critique, properly understood, does not assert that these phenomena could not or definitely do not exist. Indeed, proponents of the sensorimotor theory in some places actively assert that representations and filling-in do play a restricted role in perception. The ambition, as we saw, is to unseat the aprioristic commitment to the idea that perception depends on a bridge locus consisting of a symbolic or pictorial representation, or as O'Regan (1992) puts it, something like an 'inner cinema screen'. The inner screen conception is very close to a view explicitly proposed by Lehar (2003), but, more significantly, is liable to tacitly inform even theorists that do not explicitly endorse it, in particular - as O'Regan observes - those addressing specific phenomena such as filling-in and saccadic suppression. The sensorimotor theory's negative critique is thus similar in spirit to Dennett's (1991) earlier rejection of Cartesian materialism.

Notice that the sensorimotor theory's negative claims do not actively disprove the tenets of the orthodox approach to perception. This does not matter, as long as we appreciate that their ambition is to open up conceptual space for the sensorimotor approach's own theses. As a positive programme, the account draws on a number of unorthodox, ambitious and relatively detailed positive claims about perception and perceptual consciousness, although, as will become apparent, they are frequently in need of a great deal of conceptual clarification. Offering this clarification is my aim in the present chapter, which serves the dual purpose of laying out the tenets of sensorimotor theory and arguing for what I suggest are the most promising ways of understanding them. This will set things up for the remainder of the thesis, in which I defend the sensorimotor theory against some (albeit not all) of the best-known objections, and finish by outlining a contribution the sensorimotor approach could make to a longstanding question concerning perceptual consciousness of temporally-extended events.

3.1 - The World as an Outside Memory

We will see later in the chapter that the relation between Noë's sole-authored version of the sensorimotor theory and orthodox accounts in the tradition of Marr (1982) admits a specific complexity. Although Noë rejects the orthodox subpersonal model, his own account characterises perception primarily at the personal level. This means that a little care is required to see what sort of subpersonal story is needed to do justice to Noë's account, and which subpersonal posits the account invites, rules out or is agnostic about. In its early presentation (O&N), the sensorimotor theory focuses directly on the subpersonal processes that underlie perceptual consciousness, and therefore stands in even more direct rivalry to orthodox subpersonal accounts.

O'Regan (1992) observes that the human retina appears to have a disastrously poor design for retrieving information from the environment. Photoreceptors are increasingly limited in number the further they are from the centre of the retina, meaning that only a small central portion of the retinal image is highly resolved, while the resolution gets increasingly poor toward the outer perimeter. Moreover, the photoreceptors are, throughout the retina, unevenly rather than uniformly dispersed, and are completely absent at the blind spot, an area corresponding to a significant portion of the visual field. Visual perceivers do not ordinarily notice any of these defects, of course, and take themselves to be in perceptual contact with a richly detailed, high-definition visual scene. O'Regan claims that the disparity between the apparent richness of visual experience and the sparse and uneven detail available from the retinal image is a fundamental problem which has not traditionally been addressed with sufficient care. The result, he suggests, is that orthodox theories are frequently under the thrall of a caricatured view of vision in which the defects of the retina are compensated for by the construction of a detailed 3D model of the outside world, the activation of which is necessary and sufficient for perceptual consciousness. Almost no theorists explicitly endorse this conception, but O'Regan claims that it tacitly informs work on a range of specific visual phenomena, for instance saccadic suppression and filling-in. Notice that the conception is not disowned by Marr's very general claim that: "vision is the *process* of discovering from [retinal] images what is present in the world, and where it is" (Marr, 1982, p. 2).

The sensorimotor theory's alternative conception is founded on the claim that conscious visual experience does not comprise a richly detailed representation of the visual scene. A later step will be to re-describe visual phenomenology and the processes that enable it with attention to the fact that perceivers do not take visual experience to be sparsely furnished. The starting point, however, is to establish that the visual system does not process large amounts of detail all at once. O&N, pursuing this thesis, cite a range of empirical work, including work on inattention blindness (Wolfe, 1997), the best known example of which is the 'invisible gorilla' test (Simons and Chabris, 1999). The experiment features a short video of a basketball game with teams dressed respectively in black and white, each passing a separate basketball amongst themselves. Participants are shown the clip and asked to count how many times the players in white pass the ball, a task which requires close attention given the fast-moving action and the distracting influence of the team in black. Midway through the clip, a person in a gorilla costume enters the scene, walks casually through the middle of the game and makes a chest-beating motion. When asked if they had noticed anything out of the ordinary, around half the participants in Simons and Chabris' study failed to report noticing this take place. There are multiple ways to interpret this data, among which is the suggestion that participants were visually aware of the gorilla but did not attend to it, and therefore failed to remember it when questioned. However, the work supports the claim that perceivers do not experience unattended detail from the visual scene.

O&N draw a similar lesson from change blindness, a related phenomenon in which subjects fail to notice sudden changes to the visual scene. In an early study by McConkie and Zola (1979), participants were asked to read a line of text displayed on a monitor in alternating upper and lower case (aLtErNaTiNg CaSe). The experiment was set up so that changes to the text were timed to coincide with saccadic eye movements. During saccades, the text was modified so that the lower case letters became upper case and vice versa, although the wording, typeface, and other features stayed the same. Participants were asked if they noticed anything irregular, and universally failed to report noticing the changing case. This led the experimenters to conclude that 'visual information', meaning information derived directly from the retinal image, is not retained between fixations, although they note that the participants would have had to retain information of a more abstract, non-visual kind given

that they were able to comprehend the meaning of the text, which was read over multiple fixations.

Early work on change blindness (see also, O'Regan, 1984) focused specifically on the consequences of eye movements, and was not motivated by broader theses about the character of vision. Later, O'Regan's (1992) pre-cursor to the sensorimotor theory prompted O'Regan and colleagues to pay renewed attention to the phenomenon (Rensink, O'Regan and Clark, 1997; O'Regan, Rensink and Clark, 1999). Study participants were shown colour photographs on a computer display. Each photograph was alternated with a copy in which a key feature had been modified, for example with an object added or removed, or with the colour of the scenery changed. With the photos switched back and forth, participants could in the ordinary condition easily identify the change. The researchers found that inserting a brief flicker in-between the frames (Rensink et al, 1997) or superimposing an assortment of shapes in a 'mud splash' pattern at the moment of transition (O'Regan et al., 1999) made participants far less likely to identify the change. To explain this, they suggested that perceivers notice visual changes by having their attention drawn to sudden changes in retinal input (known as transients). This means that if the modified feature of the photograph, a local transient, is replaced by a global transient, like the flicker, or supplemented by several competing local transients, like the mud splash, this compromises perceivers' ability to detect the change.

There are various ways this observation might be incorporated into a theory in which vision also depends on the activation of an inner model. For example, the visual system may require transients to update the inner model in line with the changes. Alternatively, the inner model may be sensitive to changes in the visual scene whether or not they are signalled by transients, while transients are required to draw perceivers' attention to these changes. However, the change blindness findings make a uniquely good fit with the claim that there is no inner model, since the finding that transients are a necessary condition for detecting change holds open the more straightforward possibility that transients are sufficient for detecting change even in the absence of an inner model.

O'Regan (1992) offers an alternative framework for understanding vision, the idea of the 'world as an outside memory'. The conception is partly based on an analogy between seeing and remembering. In his examples, you might be reminiscing about your grandmother,

or attempting to remember what you had for breakfast. Long-term memories are not always consciously in mind, but typically come to awareness after being deliberately sought out. When you probe your long-term memory, past episodes do not arrive fully-formed in all their detail. Instead, you engage in a process of asking yourself questions, and answering them, and this way accessing facts about the past in a piecemeal way.

According to the outside memory view, vision works in a similar way, although the role of the memory store is played by the outside world and not the brain. The perceiver makes use of shifts in attention along with bodily movements, especially saccades, to probe the visual environment. This probing involves accessing visual detail serially, a little at a time, and in accordance with the perceiver's or visual system's present interests, i.e. the facts they want to find out and the locations they want to explore. Rather than subserving the construction of an inner model, this activity of temporally-extended probing simply takes the inner model's place.⁹

Another key analogy underlying the outside memory view of vision is touch. While vision is prone to invite an intuitive commitment to the inner model conception, it is very natural to think that perceiving the world by touch is a process of temporally-extended bodily probing rather than model-building. Noë (2004) illustrates this point with the example of a blind person who uses a cane to sample the environment one detail at a time, while making use of an accompanying grasp of the movements required to access further detail. This is how vision works on the outside memory view.

Return, briefly, to the empirical work on inattention blindness and change blindness. The outside memory thesis fits well with the suggestion, derived from work on inattention

⁹ The outside memory thesis says that perception is constituted by temporally-extended probing, and therefore might imply that the substrate of perceptual consciousness is wider than usually thought, comprising for example activity by motor areas of the brain or even activity by the extra-neural body and environment (see section 3.6 and Hurley, 1998). O'Regan (1992) stops short of claiming this, and does not offer any explicit stance on the metaphysical question of how far the substrate of consciousness extends. His early paper actively endorses the view that internal representation plays a role in vision, which might imply that the deployment of the internal representations is sufficient for perceptual consciousness without calling on any strictly necessary extra-neural activity. O'Regan should more realistically be taken as endorsing the milder claim that temporally-extended bodily probing plays an essential role in any explanation of vision.

blindness, that perceivers only experience those features that they attend to. It also predicts that visual systems rely on transients to detect change. The thesis claims that a transient occurring at the periphery of the retinal image causes the perceiver to make appropriate bodily movements to bring the object from which the transient originated into better view. Notice that the retinal image does not need to be highly or uniformly resolved to feature transients, and using transients this way could allow the visual system to detect important events at the periphery of the visual field without using an inner model. With this hypothesis, O'Regan (1992) accurately predicted the outcome of the later work on change blindness, which confirmed that change detection frequently depends on transients.

The outside memory view resonates with Brooks' (1991) suggestion, from situated robotics, that the world need not be internally modelled because it can serve as its own model. In a changing environment, it makes good sense to access detail only when required, since a notional inner model would either need to be constantly updated, which is inefficient, or risk representing the environment inaccurately. The view also observes that the visual system need not compensate for the instability of the retinal image by constructing a stable representation of the outside world. Instead, the visual system simply assumes that the outside world is a persistent and stable entity, and takes advantage of this by accessing detail directly from the environment when required, in a piecemeal and discontinuous way.

Perceivers must, of course, somehow grasp how presently accessed detail forms part of the larger visual scene. O'Regan (1992) claims that retinal stimulation is accompanied by a "non-metric awareness" (p. 474) of the locations of specific objects in the visual scene, which incorporates an understanding of the eye-movements or other bodily movements needed to access more detailed information about them via the retina. O'Regan claims that you may have this kind of awareness even of objects you have your back to, implying that the awareness is neither pictorial in its phenomenal character, nor properly-speaking visual at all. He suggests that it is enabled by neural representations which are symbolic rather than iconic, and sparse in detail. The representations do not themselves realise visual awareness, but help enable the temporally-extended probing that does.

The idea resembles the 'deictic' representation advocated by Ballard and colleagues (Ballard, Hayhoe, Pook and Rao, 1997) in which the visual system does not represent a

detailed environment, but represents sparse facts about the agent's bodily relation to objects in the visual scene. Note that the sensorimotor theory, as it emerges in later work (especially O&N), endorses the idea that perception is a process of temporally-extended exploration assisted by an implicit understanding of the movements required to bring objects into view, but does not commit itself to a particular account of the subpersonal representations that enable this, nor unequivocally to the existence of representations at all.

Return now to the apparent disparity between visual consciousness, which perceivers tend to think presents a richly detailed visual scene, and visual processing, which evidence suggests does not feature detailed representation. The apparent mismatch is captured by the work on inattentional blindness. Subjects in inattentional blindness experiments like the invisible gorilla test typically express surprise when they have the events they missed pointed out to them, indicating that they have certain assumptions about the character of their own experience (Dennett, 2001). This has led theorists like Blackmore (2002) and Dennett (2002) to propose that visual perceivers are in the grip of a 'Grand Illusion' in which they believe that they see a great deal more detail than they actually do. O'Regan (1992) claims that the apparent experience of richness is explained, in part, by the expansiveness of the retinal image, which in spite of its defects, incorporates a relatively large visual field.

This claim is not featured in the sensorimotor theory proper, but O&N endorse O'Regan's further point, which is that perceivers experience richness because they have access to missing detail when required. O&N compare this to the fridge-door effect, in which the refrigerator light seems like it is always on because it is always on when you open the door to look. This is compatible with the claim that perceivers are indeed beset by the Grand Illusion. We will see in section 3.4 that Noë makes a slightly different point, namely that the experience of detail is not an illusion needing to be explained by the fact that perceivers have access to missing detail, but a genuine experience of detail, constituted by perceivers' knowing that they have access to missing detail and implicitly knowing the precise movements that would bring that detail into view.

3.2 - Phenomenal Qualities: The Sensorimotor Contingency Account

A phenomenal quality, such as the look of yellow, the smell of a rose, or the sound of a bell, is a property of perceptual consciousness. In particular, it characterises consciousness from the first-person perspective of the conscious subject. Although we could, as a third party, describe some properties of a person, or their brain, and say we have identified the locus of some specific kind of consciousness, it is an open question whether we could ever thereby be characterising the phenomenal qualities that comprise their conscious experience. Some philosophers, most notably Jackson (1982) and Chalmers (1996), have argued that there is no way, even in principle, to characterise phenomenal qualities using physicalistic vocabulary, and hence that the qualities are not physical. Phenomenal qualities are often referred to as 'qualia', frequently though not exclusively to denote the conception of them under which they are irreducible and non-physical.

3.2.1 - The "Qualia" Problem

Jackson (1982) famously argued that qualia are non-physical using his 'Mary's room' thought experiment. The scenario features a scientist called Mary who is confined from birth to a room where she is deprived of visual stimulation in any colour other than black and white. Jackson asks us to imagine that from inside the room Mary comes to learn literally all the physical facts there are to be known, among which are all the facts that obtain about the physical events that occur when we experience colour. One day, Mary is released from the room and becomes directly acquainted for the first time with a red object. Jackson's position, at the time he formulated the thought experiment, was that Mary now learns a new fact, namely what it is like to see red, and that since she already knew all the physical facts she must therefore have learnt a non-physical fact. From this, he concluded that qualia are not physical. More recently, Chalmers (1996) has used a thought experiment to argue for the broader conclusion that phenomenal consciousness in general is not physical. He suggests that there is a conceivable and therefore, as he takes it, metaphysically possible world which is physically identical in every respect to our own, containing creatures ('zombies') that are physically identical to ourselves but lack qualia entirely. He argues that since the physical properties

instantiated in this world could be instantiated without any attendant phenomenal consciousness, a physical property cannot be identical to a phenomenal quality.

Anti-physicalist arguments like these play quite an important background role even in debates among proponents of physicalism about phenomenal character. Physicalists might deny, for one reason or another, that the premises Jackson and Chalmers use to show that qualia are non-physical are well supported in the first place. For example, some deny that there is any reason to accept Chalmers' premise that whatever is conceivable is possible (see, Gendler and Hawthorne, 2002). Others argue that given the present state of physics, we have no way to tell whether Mary, who knows physical facts as yet unknown to people in the real world, would acquire new knowledge after leaving the black and white room (Montero, 1999). In spite of arguments like these, physicalism is strategically best placed to overcome the intuitions underlying anti-physicalist thought experiments if it can provide compelling positive reasons for thinking one physical property or another is associated, necessarily, with a given phenomenal quality. For this reason, an account of phenomenal qualities has an advantage if it shows better promise than its rivals at overcoming anti-physicalist intuitions, even in a debate conducted exclusively between physicalists.

3.2.2 - Qualia and Neural-Correlates of Consciousness

The threat posed to naturalism by the suggestion that phenomenal qualities are mysterious or even non-physical is a notable and sometimes explicit background presence in the sensorimotor account. O&N begin their canonical early statement of the sensorimotor theory by criticising, in particular, attempts that have been made in neuroscience to account for consciousness by appeal to types of brain activity. They cite a number of neuroscientific accounts that have attempted to explain the existence of consciousness by appeal to facts about the neural activity it is correlated with (neural-correlates of consciousness or NCCs). These facts include, for example, the quantum properties of microtubules (Penrose, 1994), and in another case coherent oscillations within a 40 - 70 Hz frequency range, i.e. potential energy that is repeatedly displaced and restored to its original state at a constant speed of 40 to 70 times per second (Llinás and Ribary, 1993). O&N also target attempts to account for the

differences in phenomenal character between sense modalities, for example vision and audition, by appeal to the distinct neural pathways and cortical regions that are active respectively when we see and hear, a view for which they cite Müller (1838) as the historical precedent, and which they claim most present day scientists continue to be satisfied with. O&N object that all these approaches suffer from the same defect, namely that they fail to give any principled account of *why* the entities, areas or types of activity are correlated with consciousness in general, or with one phenomenal character rather than another. They argue that as long as we persist in appealing to neural properties of this sort an explanatory gap will persist, because there will always remain an unanswered ‘why?’ question.

Even if we set aside worries about anti-physicalism, accounts of the kinds described above have the drawback that they explain less than do accounts that offer a more principled explanation for why one variety of physical state rather than another yields consciousness. O’Regan and Noë’s point is even more obviously fitting when it comes to explaining degrees of difference and similarity of phenomenal character within and between sense modalities, considering that even if you believed that the relation between consciousness in general and the physical properties that instantiate it will never be intelligible to us (see, e.g., McGinn, 1991), it would still be reasonable to hope that a physicalistic account can be given which will intelligibly and non-arbitrarily explain why specific kinds of experience have one phenomenal quality rather than another. The explanatory limitations O&N identify are even more pronounced in light of the threat posed by the anti-physicalist arguments offered by philosophers like Jackson and Chalmers, since the less intelligible the proffered relation between consciousness and the physical, the greater the motivation remaining to believe that physicalism is false.

O&N do not address more promising variants of the NCCs approach. They observe that we cannot read off the properties of consciousness from the topographical quality of cortical maps that are organised in ways that replicate the retinal image, since the retinal images, being dual in number, obscured at the blind spot, and defective in various other ways, do not correlate with what we take to be the felt quality of visual consciousness. O&N do not directly address the possibility that the character of consciousness could be explained by appeal to other topological or topographical isomorphisms, nor the more promising

possibility that phenomenal character could be explained purely by appeal to the content and functional role of symbolic representations. Pessoa et al. (1998) and Noë and Thompson (2004), however, give arguments against these proposals, as I discussed already in chapter 2.

3.2.3 - The Sensorimotor Contingency Account

O&N attempt to improve on these defective explanations of phenomenal character by appealing to *sensorimotor contingencies* (SMCs), the systematic ways that stimulation received to the brain from the sense organs varies in line with movement by the perceiver's body or the objects perceived. One explanatory task that SMCs carry out is accounting for the differences in phenomenal character that hold between sense modalities, for example vision, audition and touch. As an example, O&N propose that you are facing a horizontal straight line. If you face the straight line head-on so that it is perfectly aligned with the equators of your eyes, the image projected onto your retina when flattened appears as a straight line. However, if you move your head or eyes so that the line meets your eyes above or below the equator, the eyes' curvature means that the image projected onto the retina will appear when flattened as an arc. Moving your eyes horizontally, by contrast, leaves the projection of a straight line intact. The appearance of an object also alters when projected at varying locations on the retina in line with the varying character of the photoreceptor cells. The further photoreceptors are from the centre of the retina, the greater their size, but the lower their density. Light meeting different areas of the retina is also met by a differing selection of photoreceptor types. Rods are more sensitive than cones. Cones require more light to become active, and activate, depending on their type, in the presence of differing wavelengths of light. In one part of the retina, the blind spot, there are no photoreceptors at all.

O&N call the patterns of sensorimotor contingency modulated by characteristics of the sense organs *apparatus-related SMCs*. They contrast the apparatus-related SMCs associated with the eyes with those associated with the ears, observing that sense input from the ears is not affected by blinks or eye movements, but does obey its own distinctive patterns of sensorimotor contingency. Tilting your head toward a sound source, in their example, increases the amplitude of the signal but leaves the frequency unchanged. The crucial point

these examples are supposed to illustrate is that vision and sound are not experienced as distinct conscious modalities because of the intrinsic character of the sense input signals or the intrinsic character of the brain regions associated with the respective senses, but solely on account of the distinct apparatus-related SMCs associated with the respective sense organs.

O&N also use SMCs to account for differences in phenomenal character within sense modalities. The range of phenomenal qualities that comprises a particular modality is explained by SMCs determined this time by properties of the objects perceived, rather than the sensory apparatus. These are *object-related SMCs*. Although it is intended to extend to other modalities, the account focuses on vision. Within vision, the most straightforward example of SMCs concerns shape. The sensory stimulation your brain receives from a square-shaped object changes in line with movement by the perceiver or object in a systematic way that differs from the movement-related dependence that holds when you look at a circular object. Unlike colour, the conscious experience of shape does not have any appearance of being irreducibly subjective; simply knowing that a square is a closed shape with exactly four equal sides is in principle sufficient for knowing what a square looks like. The advantage of the sensorimotor account, when it comes to shape, is that it shows how we come to experience shapes as they are, even though they are distorted when they appear on the retina. The account of shape experience is also important because it provides a partial analogue of colour experience.

Colour experience is also accounted for by the systematic dependence between movement and differing sensory stimulation, in particular the activation of differing sets of retinal cones (O&N). The analogy with shape is complicated, however, by the need to account for how the retinal stimulation you receive from an object does not just vary as your spatial alignment with it changes, but also in line with changes in ambient lighting. This point is made more explicit in the longer solo work by Noë (2004) and O'Regan (2011). We therefore need to appeal to three-way patterns of counterfactual dependence between movement, changes in ambient lighting conditions and the stimulation received by the retina. In line with the observations made about colour by O'Regan (2011) - which he points out do not constitute a final or definitive sensorimotor account of colour - this three way relation can be reduced, essentially, to a two-way relation between changes in the character of the light that

meets the object before being reflected onto the retina, and changes in the character of the retinal stimulation that follows. Note, however, that the fact that the two-way relation features fewer variables does not mean it is more appropriate than the three-way relation, which characterises the effects of movement, when it comes to explaining how perceivers actually come to be in states characterised by the relevant phenomenal qualities.

In this way, the sensorimotor account does not just appeal to properties of the light hitting the retina, but to the object's propensity in the presence of varying light sources to reflect varying wavelengths of light, and therefore stimulate the retina in systemically differing ways. O'Regan (2011) argues it is not surprising that this property of the object properly enters into an explanation of colour experience, since we are evolved mainly to respond to objects, and only occasionally interact directly with sources of light. To provide further support for the sensorimotor perspective, he cites some interesting work conducted at his lab. Philipona and O'Regan (2006) looked at changes in how photoreceptors are stimulated in the presence of a range of surfaces and lighting conditions. There are three kinds of cone photoreceptor, responding respectively to light of short, medium and long wavelengths, and each kind can be activated independently of the other two. This means that to describe cone activation in the presence of a given surface, in a given lighting condition, you need to specify a value for each of the three photoreceptors. On this basis, to specify how these change in different lighting conditions, you usually have to specify a value for three variables in each lighting condition.

The study found, however, that for surfaces with certain reflectance profiles, the explanation can be simplified, since the activations of the three photoreceptors do not vary independently in line with changes to lighting. In the case of red, yellow, blue and green surfaces, it was possible to account for changes in retinal stimulation across varied lighting conditions by specifying values in each lighting condition for only one or two variables. Philipona and O'Regan observe that this correlates with anthropological data showing that red, yellow, blue and green are the four most universally adopted colour classifications (Berlin and Kay, 1969). This provides further support, they suggest, for the idea that colour experience does not just depend on how the retina responds to reflected light of different wavelengths, but on how that responsiveness varies given changes to lighting; notice that by

just attending to the way that light of given wavelengths stimulates the retina, we would lack this explanation for the anthropological data. Since the ambient light reflected from the object to the retina varies in line with movements by the perceiver or object, the colour experience depends on a relation between movement and sense input, although to make this relation intelligible we have to also understand how lighting conditions change in line with these movements.

There is some promise, then, that the sensorimotor account of phenomenal quality can account for the experienced feel associated with distinct modalities, and within the visual modality for experienced colour and shape, solely by describing extra-neural characteristics of the perceiver's relation to the environment. This has a methodological benefit, since outward interactions are easier to observe than neural states. More importantly, it provides a more principled explanation than appealing to brain states alone. In part, this is because it appeals to patterns that are self-evidently and necessarily present in the activity of embodied perceivers. It is also because of the systematic framework it provides for explaining degrees of similarity and difference between phenomenal qualities. Noë (2004) argues that the account therefore counts as an improvement on those that suggest qualia amount to an unexplained 'residue' remaining even after we have accounted for an agent's perceptual engagement with the environment.

3.2.4 - SMCs as Counterfactual Conditionals

SMCs, as we have seen, offer a principled way to account for the varying phenomenal character associated with distinct modalities and different percepts within a modality. In particular, they also offer in principle a way to quantify degrees of difference and similarity between qualities. More has to be said, however, about what it is for a perceiver to be in a conscious state characterised by SMCs, and about the role played by the perceiver's internal states in enabling perceptual consciousness. Most versions of the sensorimotor theory complete this part of the story by reference to attunement or practical knowledge pertaining to SMCs. I will examine this notion in detail in section 3.3 and further in chapter 4, but it is worth signposting now the fact that how we understand sensorimotor knowledge or

attunement will be affected by how we resolve a key ambiguity in the notion of sensorimotor contingency.

As they are usually characterised, SMCs are appropriately described by counterfactual conditionals of the form ‘if the subject moved like this, their sense input would change like this’ and ‘if objects in the environment moved like this, the perceiver’s sense input would change like this’. This construction is offered by Noë (2004) and is also frequently in evidence throughout O&N, for example where the authors say: “Seeing an object consists in precisely the knowledge of the relevant sensorimotor contingencies - that is, in being able to exercise one’s mastery of the fact that if, among other things, you make an eye movement, the stimulus will change in the particular way typical of what happens when you move your eyes” (p. 968). On another construction, SMCs comprise the patterns of dependence between sense input and motor output that actually, rather than counterfactually, characterise the perceiver’s ongoing sensorimotor interactions. This alternative notion is found in the summary of the sensorimotor approach offered by Burhmann, Di Paolo and Barandiaran, who claim that: “According to the sensorimotor approach, perception is a form of embodied know-how, *constituted by* [emphasis added] lawful regularities in the sensorimotor flow or in sensorimotor contingencies in an active and situated agent” (2013). Here, the know-how is not knowledge about SMCs, it is knowledge which when exercised helps *constitute* the SMCs. This notion is sometimes evident in O&N, where the authors claim that exercising sensorimotor mastery consists, in the case of a guided missile, in its exercising “action recipes” (p. 945), which map certain sensory stimuli to certain motor outputs. The action recipes, the authors at one point claim, allow the missile to “*make* [emphasis added] lawful changes in sensory stimulation” (p. 945). The apparent role of the action recipes, in this case, is not to constitute knowledge *about* SMCs, but to allow patterns of sensorimotor contingency to be instantiated.

Burhmann et al. offer an explicit account of some distinct notions of SMCs in play in sensorimotor theory, which they propose can be modelled as dynamical systems. The *sensorimotor environment*, as they call it, comprises the complete set of counterfactual conditional-styled SMCs governing a perceiver’s bodily relation to its environment. This is to say that it describes the changes in sense input that would occur if the perceiver made

particular movements. Characterising the sensorimotor environment therefore only depends on the outward characteristics of the perceiver's body and the environment and not on the perceiver's brain. The authors characterise the perceiver's repertoire of actual bodily engagements within a given sensorimotor environment through what they call a *sensorimotor habitat*, which comprises the movements the perceiver makes in line with given sense inputs, and the changes in sense input that follow. As the authors note, the sensorimotor habitat is not just determined by characteristics of the extra-neural body and environment, but also by the agent's internal states.

As we have seen, the original statements of the sensorimotor theory are explicit, although only some of the time, that they mean SMCs to characterise the sensorimotor environment rather than the sensorimotor habitat. This understanding is also evident in the 'radical enactivist' versions of the sensorimotor account offered by Hutto and Myin (2012) and Myin and Degenaar (2014), which reject O&N's appeal to 'knowledge' but agree that perceivers must be 'attuned' to SMCs that exist, presumably, prior to the perceiver acquiring a sensitivity to them. The sensorimotor habitat is likely to prove useful as a way of characterising the perceiver's perceptual attunement to some subset of the SMCs in the sensorimotor environment. However, the sensorimotor habitat conception of SMCs also obscures some of the insights offered by the original statements of the sensorimotor theory. One thing it does is obscure a distinction between distinct directions of fit. SMCs, in their original formulation, describe the changes in sense input that follow changes in motor output, and never the changes in motor output that follow changes in sense input, and we could therefore think of them as having a motor output to sense input direction of fit. Even if we suppose, as if plausible, that the only SMCs relevant to an account of perceptual consciousness are those sometimes instantiated by the perceiver's actual sensorimotor engagements, it would still be true, on the original statements of the sensorimotor theory (O&N; Noë, 2004; O'Regan, 2011) that SMCs by definition have just one direction of fit, while the sensorimotor habitat describes a systematic dependence running bidirectionally between motor output and sense input.

It may turn out that the purported distinction between 'input' and 'output' is arbitrary rather than reflecting a genuine functional distinction, and therefore that the

distinction between directions of fit does not finally matter either. Nonetheless, the sensorimotor theory in its original form is explicit in individuating phenomenal qualities by appeal to SMCs with a motor output to sense input direction of fit. This is compatible with a view where inputs and outputs are instrumentally useful but dispensable parameters for modelling the workings of the wider loop of activity comprising brain, body and environment. It is also compatible with a view under which the input-output distinction is genuine. So as not to pre-judge this question, and others concerning the role played by internal states in the sensorimotor account, I restrict 'sensorimotor contingencies' (SMCs) to describing aspects of the sensorimotor environment. This of course entails that SMCs fail to explain anything about perceptual consciousness unless we say more about the role played by the perceiver and their internal states, and I examine the existing proposals about this later in the chapter. For now, I will continue by outlining the different, perhaps compatible, characterisation of sensorimotor contingency offered by Noë's (2004) solo account.

3.2.5 - Perspectival Properties and Objective Properties

Noë (2004) supplements the SMCs described by O&N with two further kinds of property, which can be thought of as new kinds of SMCs. Instead of focusing on how neural inputs from the sense organs change in line with movement, he focuses on the varying ways that objects appear to conscious perceivers. In particular, he introduces the notion of *P-property* (short for perspectival-property), an idea which is most straightforwardly explained by means of an example. A dinner plate has the 'objective' property of being circular. Viewed at an angle, however, it has the appearance of an ellipse, which gets wider or narrower as the subject faces it at different angles. The elliptical aspect is a P-property, in this case a P-shape. To experience the plate as being objectively circular, the perceiver must be sensitive to the ways that the P-shape would vary in line with movements by the object or perceiver. Thus objective properties can be characterised by reference to an amended notion of sensorimotor contingencies (SMCs). To experience the P-shape, the perceiver must understand their spatial alignment to the object in an egocentric, bodily way. Noë gives examples like knowing which bodily movements would result in your facing the object head-on, and having the ability to

indicate, using a thumb and forefinger, the shape and size of the space taken up by an object in your visual field. Thus P-property experience depends, too, on appreciating the results of possible movements, and therefore, in a way, on sensitivity to a further variety of SMCs. Noë does not say how this account relates to the account of shape perception found in O&N. We can legitimately take it that the accounts run in parallel, however, given that conscious perceivers must somehow be sensitive to SMCs styled along the lines of the earlier account in order to experience P-shapes and objective shapes.

Colour also admits a distinction between P-property and objective property. In Noë's example, you have a wall painted in a solid white. Its being a particular shade of white is an objective property. However, at a given time of day, the wall might be dappled with sunlight in one area, making parts of it appear an especially bright white, and heavily shaded in others, making it appear dark grey. This merely apparent colour is a P-property. To characterise a P-colour, we do not need to explain how retinal stimulation varies with changes to illumination, merely how the light reflected from the object stimulates the retina, and how that stimulation would change if the light met with a differing selection of photoreceptors. Here, Noë does say something explicit about the connection between his present approach and O&N. He suggests that the experience of apparent colour depends on a subpersonal process and that the sensorimotor contingencies involved in apparent colour are inscrutable to conscious perceivers, implying they are SMCs of the sort described by the earlier account rather than conscious SMCs. We experience objective colour - for instance, the colour the wall is actually painted in - by consciously exercising our practical understanding of how the apparent colour would change given changes, amongst other things, in illumination.

This parallels the suggestion given by Philipona and O'Regan (2006), although it is interesting that Noë's account of colour supplements this account by introducing this distinction between conscious and subdoxastic components of colour experience, which track the distinction between P-colour and objective colour. Noë, like O'Regan (2011), admits a final sensorimotor account of colour experience will involve further considerations still. He introduces one further proposal, namely that the objective colour we experience an object to have depends on the objective colour we take objects in its background to have. Experienced objective colour is determined, therefore, by various relations that objects stand in to each

other as well as the relation they stand in to the perceiver. Although the sensorimotor theory does not claim to offer a complete or final account of colour experience, both Noë and O'Regan make explicit their commitment to the view that colour experience is explained by patterns of movement-dependence, whether those patterns apply only to the perceiver and the object, or between the object and other features of the outside environment. Noë claims, moreover, that the capacities we use to discriminate the properties are themselves sensorimotor, a point I return to in section 3.3 and again in chapter 4.

Since the SMCs featured in Noë's account are relations between movement and, for the most part, features of conscious experience, Noë's account is committed to a distinctively controversial claim about visual phenomenology. Martin (2008) objects that the account misdescribes visual phenomenology, since we rarely experience P-properties, such as the elliptical aspect of the plate. He argues that experiencing a P-property involves adopting the 'painterly attitude', a deliberate effort of the sort a painter makes when attempting to faithfully render a scene with the correct perspective; and it is striking that doing this is actually quite difficult. In response, Noë (2008) clarifies his position by noting that you must experience P-properties even when you do not adopt the painterly attitude on account of the fact that you can only be visually aware of an object from some vantage point. Since you cannot view an object without viewing it from a particular perspective, experiencing its objective properties requires you to understand how your present perspective is contingent on your present spatial alignment to the object, and recognising this entails implicitly appreciating the ways the P-property would vary if you or the object moved. Noë claims in the same piece that you can only perceive P-properties themselves if you also experience the relevant objective properties, since you could not experience a P-property as a P-property if you did not understand the way it varies systematically in line with movement. This means that rather than being 'primitive', the ability to experience P-properties arises simultaneously with the ability to experience objective properties.

Noë (2008) also addresses an important point concerning the metaphysical status of P-properties. In at least one respect, P-properties are subject-dependent, since you cannot describe them without reference to a perceiver. Moreover, descriptions of P-properties can be used to describe the character of a subject's perceptual experience. Nonetheless, Noë (2008;

2012) emphasises that they should not be thought of as being akin to sensations, which are commonly understood as properties of consciousness and not properties of the outside world. He claims instead that P-properties inhere in the outside environment as relational properties borne by objects. As he puts it: “looks [...] are not 'mental intermediaries.' They are, rather, aspects of the things we see” (2008, p. 691) and a little later in the same piece: “I am realist about looks. My view is that the elliptical look of the plate from here, no less than its circularity, is a feature of the visible scene before me, a feature of the world itself” (2008, p. 693). It is worth exploring this claim in a little more detail. The claim that looks are properties of the world implies that they exist prior to and independently of the subject’s capacity to experience them. This seems odd, considering that the word ‘look’, applied adjectivally to a feature of the environment, implies the existence of a subject who sees it.

My suggestion is that we reconcile the idea that P-properties are in some important sense dependent on the subject’s perceptual abilities with the idea that they are properties of the world, and not the subject’s mind, by taking it that prior to featuring in a subject’s conscious awareness the P-properties exist as what Shoemaker (1980), following Geach (1969), calls *mere-Cambridge* properties. These are properties that can change without their object’s causal powers changing, and which therefore are trivial. To update an example offered by Shoemaker, we could, for instance, note that everything now existing has the property of being ‘such that Barack Obama is president of the USA’, while noting that this property could for most objects - a piece of fruit, say - change without the object’s causal powers changing. Shoemaker contrasts mere-Cambridge properties with *genuine* properties in order to denote that mere-Cambridge properties, although they exist, are uninteresting. We could contrast the mere-Cambridge cases with a scenario in which Barack Obama makes it the law that for as long as he is President of the USA, it will be a criminal offence to possess apples. In this event, the property of being ‘such that Barack Obama is president of the USA’ will mean that an apple has a causal power, namely the power to cause you to be arrested, where we take it that an object has a causal power to make X happen if its existence increases the probability of X taking place. In this scenario, the property goes from being a mere-Cambridge property to a being genuine property. Let us take it that P-properties are like this. An object might bear some P-colour or P-shape to a person who does not experience it, for example because they are

not looking at it or they cannot see. P-properties might even inhere in objects in relation to possible but non-existent perceivers. The properties would nonetheless exist, and inhere in the world, where this is contrasted with *in the mind of some perceiver*. But the properties only have causal powers, and therefore become *genuine* properties, if there is a conscious perceiver who experiences them.

3.2.6 - Bodiliness, Insubordinateness, Richness and Grabbiness

Sensorimotor contingencies are a cornerstone of the sensorimotor theory. As we have seen, their role is to explain why one sensory modality is qualitatively different from another and, similarly, why perceptual experiences qualitatively differ from one another within sense modalities. However, they do not address the broader question of why we have experiences that feel, qualitatively, like they are world-presenting rather than feeling like some other kind of experience or lacking an experiential quality altogether. The sensorimotor theory answers this question by appeal to four key notions: bodiliness, grabbiness (O'Regan, Myin and Noë, 2004; O&N), and, in later accounts, partial insubordinateness and richness (O'Regan, 2011). The properties, like SMCs, govern perceivers' sensorimotor activity in a structured environment, which is to say they help characterise the sensorimotor habitat. From a theoretical perspective, as well as from the perspective of the perceiver herself, they offer a principled way to distinguish neural signals related to perception from other kinds of signal.

Bodiliness is the tendency of body movements instigated by the perceiver to bring about changes to the input received from the sense organs. For instance, blinks and eye movements bring about radical changes to the nature of retinal input, head movements bring about changes to the amplitude of the auditory signal received from each ear, and so forth. The fact that the conscious perceiver has a degree of agential control over the character of the sensory stimulation received is supposed to help explain the sense of 'mineness' that accompanies a conscious perceptual experience. It is a necessary component of perceptual experience.

Insubordinateness is the opposite. It describes the tendency for perceivers only to have partial control over the quality of the sensory stimulation they receive. If they had full control,

the perceiver or their visual system would have no reason to think that the sensory stimulation originates from outside themselves. Thus, insubordinateness is also a necessary feature of perception, and helps distinguish neural activity involved in perception from other sorts of neural activity.

Richness is another feature that differentiates perceiving from thinking and imagining. The amount of information accessed in the course of a temporally-extended act of perception is much greater than the information available to thought or imagination. In O'Regan's (2011) example, you might consciously reconstruct, in an act of imagination, the appearance of your grandmother's face. But without seeing her, there are likely to be numerous small details about her appearance, e.g. the width of her eyebrows or the appearance of her hair line, that you cannot reconstruct. O'Regan observes that richness is not sufficient to distinguish perception from every other kind of bodily activity, since autonomic functions such as digestion and respiration also involve rich information about blood sugar levels, oxygen levels, and so forth. However, richness distinguishes sensory information from information playing a similar role in acts of thought or imagination.

Grabiness refers to the tendency of perceivers to shift their attention and bodily alignment in response to transients, sudden changes to sensory stimulation. The fact that perceivers detect change this way makes a vital contribution to the phenomenology of perception, namely the experience of being reliably in touch with a rich and dynamically-changing environment. Grabiness is a key feature of the idea of the world as an outside memory.

The above four properties, for the most part, characterise a perceiver's sensorimotor interactions in a way that can be observed from the outside. For instance, we do not need to know anything about brains or the subjective character of conscious experience to know that retinal stimulation is governed by bodiliness and insubordinateness. We can tell that richness exists purely by virtue of the rich detail available in the environment for pick-up by the retina. Grabiness is supposed to consist largely of bodily movements, in particular eye movements directed toward the source of the transients. This is complicated a little by the fact that grabiness is also supposed to be signalled by shifts in attention, which may not require eye-movements, and may instead depend solely on some kind of neural activity. With this

small complication aside, the four properties, alongside SMCs, help make the phenomenal character of perceptual experience naturalistically explicable without reference to apparently mysterious neural-correlates of consciousness.

However, perceptual experience depends on something more than SMCs and the four other principles, since it is easy to imagine objects standing in relations to the environment governed by these or similar laws without being conscious and even without demonstrating the kind of sensitivity that might justify the suggestion that they are unconscious perceivers. Perception requires that subjects are, somehow, skillfully sensitive to the fact that their relation to the environment is governed by the relevant laws. This is the purpose of sensorimotor mastery, which I discuss next.

3.3 - Sensorimotor Mastery

SMCs, on the usual understanding, are properties that could be described by counterfactual conditionals, and which pre-exist any bodily engagement with the environment, or any capacity for the same. On this basis, it is evident that to explain how a subject comes to perceive and thereby undergo consciousness characterised by the relevant phenomenal qualities, the sensorimotor account needs to show how the subject is, in a relevant way, sensitive to those SMCs. The sensitivity in question is said, in various places, to involve sensorimotor ‘skill’; ‘attunement’ to SMCs; or ‘practical knowledge’, ‘understanding’ or ‘tracking’ of SMCs. Some of the time, the variation in terminology is purely stylistic, but at other times the choice of term reflects some quite different ways of understanding the capacities in question. To keep things as clear as possible, I am going to introduce my own convention and use ‘sensorimotor mastery’ as a broad label to describe any variety of sensitivity to SMCs of alleged interest in sensorimotor theory. Under this broad heading, I will group the capacities in question under three apparently distinct types, which I will call (i) *Sensorimotor Attunement*, (ii) *Sensorimotor Tracking* and (iii) *Sensorimotor Understanding*. In appealing to these three categories I am not prejudging whether there are capacities that fall under more than one of these categories, nor whether all three deserve to make an appearance in the account.

3.3.1 - Sensorimotor Attunement

O&N suggest that over the course of your life, you are exposed to numerous SMCs, and that your brain stores information about them, categorising them under laws of sensorimotor contingency corresponding to various types of property, for example specific colours and shapes. When you are confronted with a red object, for example, your brain deploys information about the SMCs associated with red objects derived from the previous encounters. O&N do not explain what causes the brain, at a particular moment, to deploy stored information about one sensorimotor law rather than another. We can plausibly suppose that the brain does not just make use of present sense input state to detect which SMCs obtain, but also efferent copy from the motor regions, and the way both have changed in the last few moments.¹⁰ This is to suggest that the brain responds to information about the restricted set of SMCs evident in the bodily interactions that are taking place right now, and associates these SMCs with stored information about the fuller set of SMCs that usually apply in similar situations. I think it plausible that this picture, or something close to it, is what O&N have in mind, although they do not say so explicitly.

The authors are emphatic that neither the availability of the stored information nor its mere ‘activation’ is sufficient for perception. Instead, they argue that perception is a kind of embodied exploration of the outside environment analogous to a dance, an activity that can never consist of an inner state alone, but also requires a variety of temporally-extended interaction with the outside environment. They say: “Seeing is not constituted by the activation of neural representations. Exactly the same neural state can underlie different experiences, just as the same body position can be part of different dances” (O&N, p. 966). This invites an understanding of sensorimotor mastery as a capacity for making a variety of systematic, goal-directed bodily movements. This capacity, which I call Sensorimotor Attunement, can also be found in Burhmann, Di Paolo and Barandiaran’s (2013) model, which invites us to model a perceiver’s repertoire of bodily engagements in a structured

¹⁰ It is not a logical necessity, however, that the neural system which enables conscious vision is always sensitive to ongoing efferent copy. It might only be sensitive to efferent copy during learning phases, as Clark (2008) suggests, after which it uses sense inputs alone to recognise patterns of sensorimotor contingency.

environment as a sensorimotor habitat.

Although O&N and Noë (2004) on one or two occasions deny that vision is exhaustively accounted for by a set of behavioural capacities, the idea that vision depends on being poised to act is recurrent throughout their accounts, for example their claim that: “Under the present theory, visual experience does not arise because an internal representation of the world is activated in some brain area. On the contrary, visual experience is a mode of activity involving practical knowledge about currently possible behaviors and associated sensory consequences. Visual experience rests on know-how, the possession of skills” (O&N, p. 846), and in Noë’s account: “It is my possession of basic sensorimotor skills (which include the abilities to move and point and the dispositions to respond by turning and ducking, and the like) that enables my experience to acquire *visual* content” (Noë, 2004, p. 90).

The idea of Sensorimotor Attunement as a behavioural capacity is evident, in particular, in the main example given by O&N, which features a guided missile. The missile’s goal is to hit its target and, while following the target, to try to avoid enemy fire by zig-zagging left and right. The missile’s relation to the target is characterised by various SMCs. When the missile turns left (on its yaw axis), the target, if it appears in its sensors, will appear further to the right-hand side, and vice versa. If the missile gets closer to the target, the target will occupy more space in its sensors. To achieve its goal, the missile’s behaviour must exhibit sensitivity to SMCs like these. O&N, who are emphatic that they do not wish to underplay the contribution made by the brain to perception, propose that the missile, like a human, makes use of stored information, which in the missile’s case comprises “action recipes” (O&N, p. 945) corresponding to various sorts of target. They emphasise, however, that the internal states do nothing until they are deployed for some practical end. In the case of the missile, they observe that the action recipes are “latent” (O&N, p. 945) until the missile begins following a relevant target.

The authors do not say exactly what an action recipe is, but evidently it comprises instructions about how to move. Presumably, it does this by mapping sense inputs to motor outputs, for example specifying that the missile should turn right if the target appears on the right hand side of its sensors. Attending to the contrasting descriptions of the stored information deployed by humans and the guided-missile when exercising Sensorimotor

Attunement, we might notice that direction of fit, which has already been an important notion in understanding SMCs, is an issue again here, although in a different form. The stored information about sensorimotor laws that O&N allege to be involved in human perception appears to have a belief-like direction of fit - it specifies how sense input would change given particular movements, and it must therefore correspond to pre-existing facts about the sensorimotor environment. Action recipes, on the other hand, appear to have a desire-like direction of fit - they specify which movements will take place given particular sense inputs, suggesting that if they are effective it is the missile's behaviour that will appropriately correspond. The fact that O&N refer interchangeably to mastery of SMCs and the possession of 'action recipes' has the potential to be puzzling if we do not unpack the notion of Sensorimotor Attunement with care, since it is not clear straight away why we should think of the belief-like and desire-like capacities as being identical. To this end, we must attend carefully to O&N's claim that the exercise of Sensorimotor Attunement is never *solely* a matter of stored information being activated by the presence of relevant stimuli, but a matter of the stored information being activated in the course of the agent's goal-directed behaviour.

O&N insist that the attunement, is "a practical, not a propositional form of knowledge" (p. 944). Given their characterisation of sensorimotor knowledge as knowledge *that* particular SMCs obtain, we might be drawn to suppose that despite not being propositional, it nonetheless has a belief-like direction of fit. On this basis, their insistence that it must be exercised might reflect nothing more than the commonplace distinction between possession and exercise that applies to any kind of knowledge. For example, your knowledge that Edinburgh is the capital of Scotland only features in your awareness if you are exercising it, perhaps because you are presently thinking about travelling to the capital of Scotland. This offers one reason for thinking sensorimotor knowledge or attunement must be exercised. O&N endorse a stronger claim, however, when they say: "the notion of being tuned, or having [sensorimotor] mastery, *only makes sense* [emphasis added] within the context of the behaviour and purpose of the system or individual in its habitual setting [...] if exactly the same system was being used for a different purpose [...] it might well be necessary for the system to have a different behavior [in order to exercise sensorimotor mastery]" (p. 943). O&N cannot simply be claiming here that a system *ought* to behave differently if its purpose

is different, since this would be trivial. Instead, the passage suggests that the attunement to SMCs is only something that exists by virtue of the agent's purpose and behavioural repertoire.

If Sensorimotor Attunement can only be made intelligible by appeal to a capacity for goal-directed behaviour, as the above quote appears to suggest, it depends on the existence of a capacity with a desire-like direction of fit. In this case, it is reasonable to ask where the belief-like attunement to SMCs comes from. One way to unpack O&N's point is by supposing that we specify which SMCs the perceiver is attuned to by interpreting its behaviour in light of the SMCs that now obtain, along with what we take to be the perceiver's goal state. This might involve something analogous to Dennett's (1987) proposal that we properly ascribe beliefs and desires to a subject on the basis of their explanatory value in making sense of the behaviour. In the case of Sensorimotor Attunement, we first might notice that the missile is disposed to yaw to the right when target appears in the right hand side of its sensors. Next, we might suppose that the missile's purpose is to keep the target aligned in the middle of its sensors; this ascription of purpose could either be a theoretically useful assumption, as with Dennett's model, or we might actually know that this is its purpose because we know what the missile's designer intended. With this purpose in mind, we can make sense of the missile's behaviour by ascribing to the missile attunement to the relevant SMCs, in this case the fact that turning right will result in the target appearing in the middle of its sensors. The 'action recipes' may work by mapping possible sense inputs to appropriate motor outputs, while the action recipe selected may co-vary with different sets of SMCs. It is the interpretive stance, however, that on this proposal admits the dual directions of fit.

O&N do not go so far as to overtly insist that 'purpose' is necessary to specify which SMCs the perceiver is attuned to, either by using the interpretive strategy sketched above or any other strategy. Indeed, neither teleology nor having a proximal goal are given a prominent theoretical role either in the remainder of the paper or in other statements of the sensorimotor account. It is probable, then, that O&N merely mean to suggest that having a purpose is necessary for Sensorimotor Attunement to exist in general. On this basis, their view can be taken as the more austere claim that the perceiver is attuned to a set of SMCs if it is able to respond with a distinctive set of movements when those SMCs obtain, and different

movements when other SMCs obtain, with the added condition that for this sensitivity to constitute *attunement* (or knowledge, skill, etc.) it must somehow help the perceiver fulfil a purpose or achieve a goal. This latter proviso seems right, considering, for example, that when a lake reflects the night sky it demonstrates a sensitivity to the night sky, but does not thereby manifest any skill.

It seems viable, then, that Sensorimotor Attunement could be regarded as nothing more than a behavioural capacity of the sort I have just described. O&N appear to disavow such a view, however, when they make an effort to rebut charges that their theory amounts to a form of behaviourism. In particular, they say: “[I]t is *not* the case that mental states (experiential states) are logical constructions of actual and possible behaviour states” (2001b, p. 1014), and later: “Knowledge of the laws of sensorimotor contingency themselves must surely be represented” (2001b, p. 1017). We could attempt to do justice to this with a representationalist gloss in which Sensorimotor Attunement involves the deployment of representations with both belief-like and desire-like directions of fit, i.e. representations of the SMCs that now obtain and the behaviour that should follow. The notion of a neural representation with both directions of fit is found in the literature on *action-oriented representation* (e.g., Wheeler, 2005; Clark, 1997), and if we wanted to give explanatory pride of place to the neural state that realises the behavioural capacity, it would on this view be the content of the representations that determines the SMCs to which the perceiver is now attuned. However, it is not clear that this fits with O&N’s claim that: “Seeing is not constituted by the activation of neural representations. Exactly the same neural state can underlie different experiences, just as the same body position can be part of different dances” (p. 966).

I think the best way to understand Sensorimotor Attunement is to suppose that while other varieties of sensorimotor mastery involved in perceptual consciousness may draw on something other than a behavioural capacity, Sensorimotor Attunement is, in fact, a logical construction of actual and possible behaviour states. We can do justice to O&N’s resistance to appealing to nothing but behaviour by adding a couple of provisos to the present notion of Sensorimotor Attunement. A first is that a full explanation of the faculty must also explain, by appeal to stored information, representation, or something else, how the behavioural capacity

is realised. The present notion of Sensorimotor Attunement denies that these internal states can be explanatorily relevant if they do not explain the behavioural capacity, and further denies that the perceptual state can change without the behavioural capacities changing. Whether or not the notion amounts on this view to a form of behaviourism is open to debate, but it certainly has the resources to refute the suggestion that because it is a form of behaviourism it ignores or denies the important role played by the brain.

The second proviso is that the behavioural capacities constituting Sensorimotor Attunement cannot be accounted for by a simple mapping from varieties of sense input to varieties of motor output, but only by a mapping from laws of sensorimotor contingency to varieties of motor output. O&N suggest this entails that because sensorimotor laws vary in line with differences in the body's morphology, minute differences in embodiment between one perceiver and another entails that the two perceivers must be discriminating a different range of sensorimotor laws. I will argue later that the sensorimotor theory need not be committed to the view that perceptual states are maximally sensitive to the details of embodiment, although it ought to say that differences in embodiment between two perceivers necessitates that they are making a different range of perceptual discriminations when the differences are big enough.

Going further beyond the explicit claims made by O&N, I am going to add one further criterion for establishing that a perceiver exhibits a belief-like attunement to SMCs of the sort we require. Wheeler (2005) describes work by Harvey, Husbands and Cliff (1994) featuring a robot, designed using artificial evolution, which was able to achieve a practical end by following an identical behavioural repertoire for all sense input states other than the target state. Placed in an environment containing a triangle and a rectangle affixed to a wall, the robot's task was to locate the triangle and move toward it. Its strategy was to do nothing but rotate until its sensors were stimulated in a way characteristic of the triangle, and then move forward. Faced at an angle, the rectangle's edge would appear in the robot's sensors as a diagonal line, and in cases where the rectangle's perspectival shape was sufficiently similar to the triangle's perspectival shape when faced head-on, the robot would 'mistake' the rectangle for the triangle and begin to move forward. However, once the robot had moved some way toward the shape, the rectangle would stop presenting the triangle-like aspect, causing the

robot to resume rotating and eventually head toward the triangle, which would continue to present the appropriate perspectival shape as the robot got closer.

We can observe that the robot certainly possessed a rudimentary degree of Sensorimotor Attunement. It was implicitly attuned to the fact that continuing to rotate while it had the triangle in its sights would result in its losing sight of the triangle, and to the fact that proceeding straight on toward the triangle would allow it to keep the triangle in view. More interestingly, its behaviour over a short time period also manifested a sensitivity to the fact that although the rectangle had the propensity to appear like the triangle when faced from certain positions, it did not from others. This latter sensitivity is not explained by stored information or historical facts about the SMCs that the robot had been exposed to, but by properties of the rectangle itself. Wheeler (2005), following Wheeler and Clark (1999), describes this sort of environmental contribution as *non-trivial causal spread*, denoting that the environment itself plays some explanatorily interesting causal role in the production of the behaviour.

O&N do not say anything explicit to indicate whether temporally-extended engagements demonstrating sensitivity to SMCs, where that sensitivity is accounted for by non-trivial causal spread, should be counted as instances in which the agent exercises Sensorimotor Attunement to the same SMCs. The most cohesive way to understand the sensorimotor account is to suppose that they should not. This reflects the fact that O&N actively emphasise the role played by internally stored information derived from previous encounters with the environment. Indeed, sensitivity to SMCs accounted for by non-trivial causal spread cannot make any contribution to O&N's view that when we exercise the right kind of sensorimotor mastery we are conscious of SMCs that are not capturable by the patterns of output-input dependence manifest in our present movements. For example, you can be conscious of a wall's objective colour without any changes in illumination taking place; and you can feel embedded within a richly detailed visual environment without taking in all that detail, even in the course of a temporally-extended sensorimotor engagement, a point I will discuss in more detail later in the chapter.

Since non-trivial causal spread cannot account for many of the key phenomena the sensorimotor theory is supposed to account for, I think it best to consider the non-trivial

causal spread case to demonstrate that relatively basic Sensorimotor Attunement *plus* non-trivial causal spread can result in surprisingly sophisticated patterns of sensorimotor engagement, but avoid attributing those patterns to Sensorimotor Attunement alone. Another way to put the same point is to stipulate that possession of Sensorimotor Attunement cannot depend on anything temporally-extended, even if it is exercised in the course of temporally-extended behaviour. To be in the relevant way attuned to SMCs that now obtain, a perceiver must have the ability to make an immediate change to its motor outputs demonstrating its sensitivity to those SMCs.

Sensorimotor Attunement, the way I have characterised it, reflects a notion found at least sometimes in the main statements of the sensorimotor theory. It is also, in principle, compatible with Myin and Degenaar's (2014) notion of 'attunement' to SMCs, although it is worth noting that their notion is distinct. Sensorimotor Attunement, as I have presented it here, implies a capacity for outwardly observable behaviour, with representation playing no more than an instrumental role in enabling this behaviour. Myin and Degenaar also describe attunement to SMCs as an embodied capacity, but say: "in the sense of 'embodied' at stake here, 'embodied' is used in contrast with 'representational'" (p. 1). By this, they mean to contrast their preferred account with those accounts that mention embodiment *and* representation. However, Myin and Degenaar's preferred understanding of 'embodiment' also stands in contrast with accounts that make explicit appeal to capacities for outward behaviour.

They stop short of making any such appeal to any such capacities, referring only to the agent's *historical* sensorimotor interactions. They say: "Attunement to sensorimotor contingencies then means that an organism has acquired, on the basis of a history of interactions, a sensitivity in its perception and action for the ways stimuli change with movement" (p. 5). But other than denying that it involves representation, they do not say anything about what a sensitivity to SMCs in perception is constituted by right now. Note, then, that although Sensorimotor Attunement as I have presented it does not rule out representation, it provides a positive non-representational account of what it is to possess the relevant capacity, namely the disposition to manifest different behaviour when different SMCs obtain. This means the notion I outline now has an explanatory advantage compared to Myin and Degenaar's, which is limited in that it only tells us what caused the organism to be

in its current state, and not what it is about the organism's current state that results in it being in a perceptual state.

3.3.2 - Sensorimotor Tracking

O&N note the existence of an apparently key limitation with the notion of Sensorimotor Attunement as I characterised it in the last section - namely that since the guided missile, for example, is not conscious, Sensorimotor Attunement is not sufficient for perceptual consciousness. To account for visual awareness, O&N introduce the further criterion that you must be exercising your attunement to SMCs for the purpose of "thought and planning" (p. 944). To illustrate, they give the example of driving a car while talking to a friend. While driving, you must navigate a range of obstacles while modulating your driving behaviour appropriately, yet you are not necessarily aware of everything you are doing. In this respect, you are equivalent to the autopilot of a plane. Only when you attend, for example, to the colour of the car in front, and think or talk about it, or make plans in accordance with it, do you become aware of the colour. One way to understand this requirement is to suppose that the Sensorimotor Attunement only yields perceptual consciousness when it is deployed in certain kinds of bodily movement. Another possibility is that it depends on a certain kind of neural representation; O&N do not give a particularly detailed proposal about what happens physically when thought and planning takes place. I will call sensorimotor mastery in this context Sensorimotor Tracking.

In his book *Why Red Doesn't Sound Like a Bell*, O'Regan (2011) builds on the account offered by O&N. The book explains in greater detail the relation between Sensorimotor Attunement, which by itself does not necessarily involve consciousness, and the thought and planning needed for conscious perceptual awareness. The key faculties described in O'Regan's model of perception and perceptual consciousness comprise four hierarchically-organised tiers. The bottom tier - call this Tier 1 - comprises the agent's skill-driven sensorimotor coupling with the outside environment. This type of coupling could be instantiated by a human, but also by an unconscious artifact like the guided missile.

O'Regan's account of sensorimotor mastery and SMCs at Tier 1 at times appears to

adopt the notion of SMCs in which they are constituted by patterns present in the agent's repertoire of skillful behaviour (its sensorimotor habitat), rather than the notion in which the skillful behaviour involves a behavioural sensitivity to counterfactual conditional-styled SMCs that pre-exist in its sensorimotor environment. As I have already argued, it is better to explicitly reject this understanding of SMCs, for reasons of expediency. The understanding obscures the distinction (which *may* be theoretically useful) between patterns of dependence between sense input and motor output, and those between motor output and sense input. It also obscures the distinction between SMCs which are merely manifest in this behaviour, and those to which the agent's behaviour demonstrates a skillful sensitivity. Finally, it muddies our understanding of SMCs in light of the contradictory conception on them found in Noë (2004), throughout much of O&N, and at times in O'Regan (2011), in which they are the ways that sense input *would* change given particular movements, regardless of whether those movements take place. Fortunately, it is also consistent with O'Regan's (2011) account if we regard Tier 1 as involving unconscious sensorimotor coupling drawing on Sensorimotor Attunement to SMCs as I have characterised the concepts in the last sections.

The next tier in O'Regan's hierarchy - call this Tier 2 - involves a higher-order cognitive access to SMCs governing activity at Tier 1. This is Sensorimotor Tracking, and is exercised when the subject is poised to exhibit sensitivity to sensorimotor laws in the course of "planning, decisions, or other rational behaviour" (O'Regan, 2011, p. 121). The behaviour in question falls short of self-conscious behaviour or discursive thought and appears to be quite intimately connected, still, with bodily behaviour; O'Regan offers as an example of behaviour at this tier the state of being poised to press on a brake pedal when a red traffic light is visible. It is not clear whether or not we should suppose that the only SMCs you can be sensitive to in Tier 2 activity are those you already exhibit Sensorimotor Attunement toward in your unconscious activity. However, O'Regan's account does not make this an obvious requirement. So it is more straightforward to suppose that Sensorimotor Tracking at Tier 2 involves exhibiting sensitivity to SMCs, in general, for the purpose of behaviour which is planned and answerable to rational norms. This depends on tracking or being sensitive to the SMCs governing your unconscious activity at Tier 1. But it may involve being sensitive to SMCs beyond those featured in your Tier 1 Sensorimotor Attunement, considering that some

of the patterns manifest in your sensorimotor activity at Tier 1 may not be SMCs to which this activity is itself Attuned, considering, to take one example, the proviso I made in the last section that Sensorimotor Attunement does not encompass skillful sensitivity to SMCs when this is accounted for by non-trivial causal spread.

O'Regan suggests that Tier 2 activity involves the sort of cognitive access to some feature that Block (1990a) would describe as having 'access consciousness', since it involves being poised to make use of that feature for reasoning and planning. O'Regan denies that having this first order of cognitive access is sufficient to constitute being conscious, however. To be conscious, O'Regan suggests we must have a further order of cognitive access to activity at Tier 2. This further higher-order tracking, call it Tier 3, constitutes another level of Sensorimotor Tracking. Perceptual consciousness occurs when the perceiver thinks reflectively about features of the perceptual environment corresponding to SMCs which it already tracks at Tier 2. The subject does not have access, when engaging in this kind of thought, to the individual SMCs that correspond to a particular property, but merely to categories or descriptors corresponding to sensorimotor laws, e.g. 'that shade of red'. This kind of thought is not intimately connected with bodily activity. However, in common with the activity described by the tiers below, it is always practically focused, involving active, goal-directed interrogation of the visual scene. In O'Regan's example, this involves asking and answering questions like: "Is it an 'A'?; 'Is it Times New Roman?'; 'Is it black on white?'" (2011, p. 28).

In spite of the more fully-developed character of O'Regan's (2011) account, it lacks a detailed characterisation of the neural (or neurally-realised, functional) processes that enable activity at Tiers 2 and 3. At one point O'Regan suggests his account may bear an affinity to Higher Order Thought (HOT) theories of consciousness, although he is explicit about not being firmly committed to this suggestion. HOT theories explain consciousness by appeal to higher-order access to lower order states in a manner similar to O'Regan's model. Notably, they almost always cash this suggestion out by appeal to a subpersonal explanatory level in which one part of the cognitive systems bears higher-order representations of activity in another system. If the tiered sensorimotor theory adopted a view like this, it would be committed to explaining Tier 3 by reference to a subpersonal representation of Tier 2 activity,

and Tier 2 by reference to subpersonal representations of SMCs. This sort of proposal will be particularly relevant in chapters 5 and 6, where I discuss the relation between sensorimotor theory and the dual systems model of vision.

O'Regan (2012) later asserts that the tracking carried out by Tiers 2 and 3 involves neural states that 'track' or 'register' the appropriate lower-order activity, but do not constitute representations. The neural states register the appropriate interactions by virtue of being the sorts of physical states that ordinarily occur when the appropriate interactions are taking place. O'Regan denies that these states thereby count as representations, although it is notable that some accounts of representation (e.g., Tye, 2000, or Rupert, 2011) appeal to no more than simple covariance relations of this sort to constitute representation, as I will discuss in chapter 5. As it is, O'Regan's (2012) notion of tracking appears to correspond to Myin and Degenaar's (2014) notion of attunement to sensorimotor contingencies (not to be confused with 'Sensorimotor Attunement' as I use it here) in which the agent is attuned to sensorimotor contingencies by virtue of being in some unspecified physical state while having the relevant history of sensorimotor interactions.

The notion of tracking or registering just described appears to be inadequate by itself to do justice to O'Regan's (2011) multi-tiered account, however, inasmuch as it does not explicitly specify the role these physical states perform when deployed in reasoned behaviour (at Tier 2) or reflective thought (at Tier 3). Considering these activities are supposed to be essential for enabling perceptual consciousness, it is reasonable to desire a better developed, naturalistically-respectable account of what these higher order capacities involve. One possibility is that Tier 2 and perhaps even Tier 3 capacities depend only on capacities for behaviour, for example behaviour that is suitably sophisticated, or answers to appropriate rational norms. O'Regan does not say much about this, but suggests possible sympathy toward the suggestion that we ascribe these capacities to an agent through a Dennett-style interpretive stance similar to the one I discussed in Section 3.3.1. O'Regan hints at a position of this nature in a footnote, where he declares: "The implication is that conscious access is not something in the agent. It is a way that an outside person (and consequently the agent itself) can *describe* the state of the agent" (2011, p. 125n1).

3.3.3 - Sensorimotor Understanding

Sensorimotor Understanding is the variety of sensorimotor mastery featured in Noë's (2004; 2012) sole-author versions of the sensorimotor account. I examine the capacity in greater detail in the next chapter, so for now I will restrict myself to a brief outline of its key features. In contrast with the other varieties of sensorimotor mastery, Sensorimotor Understanding pre-supposes consciousness. For this reason, Noë does not present his version of the sensorimotor account as an attempt to dissolve the hard problem of consciousness (cf., O&N; O'Regan, 2011) and Sensorimotor Understanding is instead merely intended to help characterise perception and perceptual consciousness in an already conscious perceiver. Noë (2004) describes the capacity as an implicit understanding of P-properties, the understanding of which requires an appreciation of your egocentrically specified alignment to objects in the environment, and objective properties, the understanding of which requires an appreciation of how P-properties would change in line with movement. Although he downplays the importance of the personal/subpersonal distinction, Noë's (2004) characterisation of Sensorimotor Understanding presents it primarily as a personal-level phenomenon. However, he suggests that for some kinds of experience, including P-colour experience, we do not have conscious access to our exercise of sensorimotor mastery, thereby implying that in limited cases it is firmly subpersonal.

Noë and O'Regan's respective sole-author accounts offer quite different treatments of sensorimotor mastery and there is little to suggest that Sensorimotor Tracking and Sensorimotor Understanding should necessarily be regarded as different descriptions of the same notional faculty. However, given that they both originate from the early O'Regan and Noë paper, they could, with a little finessing, be thought of as different treatments of the same phenomenon. If so, Sensorimotor Understanding could be regarded as a more philosophical exposition of Sensorimotor Tracking. The main obstacle to such a reconciliation is the difference between the properties the respective accounts claim are the objects of sensorimotor mastery, namely subdoxastic sense inputs (O&N; O'Regan, 2011) and P-properties and objective properties (Noë, 2004; 2012). Noë's (2004) account does make reference to the patterns of movement-related change governing input from the sense organs,

claiming we must master them to experience P-colour. This indicates that Noë does not intend his sole-author account to entirely supersede the early statement of the sensorimotor theory, and thereby makes a reconciliation between his and O'Regan's later solo accounts look more promising. Sensorimotor Tracking and Sensorimotor Understanding notably both lack a proper account of the physical or subpersonal mechanisms they are realised or constituted by.

Sensorimotor Understanding is uncertainly defined in a couple of key respects. One, as I just suggested, concerns the capacity's place in a scientific account of perception, in particular whether or not it depends on internal representation, and more narrowly the precise representational or non-representational mechanism that enables it. Another concerns Noë's (2004) claim that it is a kind of knowing-how, and not a kind of knowing-that, which is in tension with his firm insistence, elsewhere in the book, that Sensorimotor Understanding is not knowing how to act but knowing what would happen if you did act. I make some suggestions about how we should resolve these issues in chapter 4.

3.4 - Perceptual Presence: As-Access and In-Absence

As we have seen, the sensorimotor theory frequently criticises the idea that vision is enabled subpersonally by the neural activation of a detailed or, in particular, pictorial representation of the outside world. The outside memory thesis says that the brain does not internally model the outside environment but instead accesses information directly from the outside world. This entails that perception, considered subpersonally, depends non-trivially on bodily interaction with the outside environment, and therefore does not depend solely or primarily on internal representation.

In his sole-authored versions of the sensorimotor theory, Noë turns his attention to a personal-level analogue of the subpersonal views targeted by O&N. The 'snapshot conception' (Noë, 2004) endorses a pictorial understanding of the phenomenology or what-it-is-likeness of visual experience, suggesting that visual experience presents in pictorial fashion a large and uniformly detailed visual field. Noë (2012) draws a connection between this phenomenological claim and the logical claim that perception at the personal level involves

representing the outside world. Declaring an allegiance with relational realism (e.g., McDowell, 1994a; Campbell, 2002), he denies that conscious perception involves personal-level representation, and claims instead that it involves standing in a specific kind of relation to the outside world itself. This point is reflected by the word he uses to label the objects and properties to which we are intentionally-directed in conscious perception, which he describes as being 'present' rather than 'represented'.

Noë claims that objects and properties in the outside environment become perceptually present when we stand in a particular kind of bodily relation to them. As we have seen, that bodily relation can be specified by counterfactual conditionals describing the effects of possible movements, with the relevant counterfactual conditionals being ones the perceiver implicitly understands. Noë's (2012) relational view of intentionality generalises to encompass non-perceptual forms of thought, with perceptual presence being distinguished from other forms of presence by the fact that the perceptual relation belongs to this particular type. Crucially, the presence perceivers achieve by standing in this relation is not the simple kind of presence characteristic of the snapshot conception, but a special kind of presence which Noë calls presence-as-access.

The idea of presence-as-access is best illustrated by how it accounts for a feature of visual phenomenology called presence-in-absence. Noë's favourite example features a tomato (2004; 2012). As with any opaque, three-dimensional object, you can never see every side of a tomato at once - you see one side, while another side remains out of view. However, Noë claims that when you look at a tomato, the rear side is nonetheless perceptually present 'as absent'. This is motivated by attention to visual phenomenology. Noë takes it as a primitive that the back of the tomato has perceptual presence even though it is not visible in the way that the front of the tomato is, a claim he justifies in part by noting that it does not seem as if you are merely looking at a two-dimensional cut-out of a tomato. The perceptual presence of the back of the tomato is enabled, he suggests, by your understanding of the movements that would bring the rear of the tomato into view.

The perceptual presence of objects and properties that are out of view is a pervasive feature of perceptual experience, and Noë emphasises the distinctive ability of presence-as-access to account for it. Another central case is the perceptual presence of

occluded parts of an object, for example a cat sitting behind a fence (Noë, 2002; 2004). Although parts of the cat are occluded by the fence posts, a true characterisation of visual phenomenology will observe that a visual experience of the cat is not an experience of a cat's parts, but of a whole cat. This is accounted for, as before, by the perceiver's understanding of the movements they or the cat would have to make to bring the occluded parts into view.

Presence-as-access is also intended to account more generally for the perceptual presence of a richly detailed visual environment. Although the visual system does not process a lot of environmental detail all at once, as evidenced by the experiments on change blindness and inattention blindness, Noë (2002) suggests that visual perceivers experience the environment's rich detail as present thanks to their implicit understanding of the movements, such as eye movements, that would bring relevant detail into view. Interestingly, he denies that perceivers ordinarily suffer from any delusion about the character of their access to environmental detail (cf., Dennett, 2001). He suggests that the surprise participants in Simons and Chabris' (1999) inattention blindness experiment reported on finding out that they failed to notice the gorilla walking through the scene is explained by the implicit expectation that their attention would be directed toward the gorilla, rather than the belief that visual experience presents rich detail all at once.

Noë (2012) claims that there is no firm dividing line between objects that are perceptually present and objects that are not. He claims that a perceiver may even have quasi-perceptual experience of an object in the very distant environment, such as a friend living in a distant city. If you know where that friend lives, your thoughts about him are characterised, in part, by an understanding of the fact he is far away and an understanding of the movements you could make to reach him, such as getting on a plane. For this reason, the thought takes on, in part, a quasi-perceptual character, although Noë does not claim this counts as a proper instance of perceptual presence. He contrasts this with the case of a friend with whom you have lost touch. Because you do not know where the friend lives, or whether he is alive or dead, the phenomenology of thoughts about this friend is quite different, a difference explained at least in part by the fact that thoughts in the latter case do not draw on an understanding of your sensorimotor relation to the friend.

The most striking feature of presence-in-absence (the phenomenology) taken together

with presence-as-access (the logical claim that accompanies it) is that all perceptual experience is supposed to belong to this type. This means presence-in-absence and -as-access incorporates experiences we might be tempted to think feature a simpler kind of presence, for instance the experience of the front side of the tomato. This is a counter-intuitive and, on a certain reading, highly problematic claim, so it is worth noting that Noë is explicit on this point. He says: “*all* perceptual presence is presence as absence [...] perceptual presence is availability; the world shows up for us as accessible” (Noë, 2012, pp. 95-96). Instead of imagining that the tomato’s facing side has presence simpliciter while the non-visible parts are presented as absent but accessible, Noë is explicit that the perceptual presence of the facing and the rear side both belong to the latter kind. This is supposed to be evident from introspectable features of visual phenomenology. In a description duplicated with slight modification from a passage originally found in his earlier book, he says:

[S]top and look again, for example, at the face of the tomato. You can’t comprehend the whole of it all at once in your visual consciousness. You focus on the color now, but in doing so, you fail to pay attention to the shape, or to the variations in brightness across the surface, but only at the price of ignoring the rest of it. You can no more achieve perceptual consciousness of the whole aspect of the tomato’s front side all at once than you can see the tomato from every side all at once [...]. In a way, for perception, everything is hidden. Nothing is given. (Noë, 2012, p. 19)

The observation that we cannot fixate on atomic features of the visual environment is compelling, but the suggestion that everything is hidden is more obscure and can be glossed in a number of different ways, as we will see.

In his earlier account, Noë (2004) gives a similar description of visual phenomenology and suggests it falls out of a particular understanding of the physical process that underlies perceptual consciousness. To this effect, he offers an analogy between visual perception and a computer accessing an online version of a newspaper via the World Wide Web. Rather than downloading the entire issue at once, the computer downloads the articles one at a time, accessing them only when called upon by the user. Downloading the issue in a piecemeal way has certain advantages. It limits the burden placed on the computer’s bandwidth and processing power, and it means that if an article is amended, the user can rely on seeing the most up-to-date version. Notice that this is similar to perception, understood in accordance

with the ‘outside memory’ thesis that perceivers use eye movements to access missing detail, when required, directly from the outside world. Importantly, it makes no practical difference to the person reading the newspaper that the computer downloads the articles one-by-one rather than all at once, since they have immediate access to the articles as they request them. As Noë puts it, the computer and its user therefore have ‘virtual’ access to the entire newspaper. This is comparable to conscious vision, in which he alleges not only that perceivers access detail in the way suggested by the outside memory thesis, but also that they have a phenomenological sense of that detail’s presence by virtue of knowing how to access it when required. Foreshadowing the claim that *all* perceptual presence is presence-as-access, Noë claims that conscious vision is in part disanalogous to the computer case, because rather than factoring into an offline and online component, perceptual presence is “virtual *all the way in*” (2004, p. 134).

The computer analogy may invite a problematic understanding of the claim that in perception, everything is hidden. While the computer ‘knows’ how to access the New York Times, it does not anticipate, in advance, what the contents of the articles will be. Noë’s descriptions of presence-as-access frequently imply that conscious accessibility works in a similar way. On a certain reading, this suggests that subjects never consciously perceive the environment, and that all they are conscious of is a disposition to make exploratory bodily movements. The claim that perceptual consciousness consists of no more than this is suggestive of an outright denial of perceptual consciousness. Noë (2004) explicitly disavows the suggestion that denying perceptual consciousness is a consequence of the account. If it was a consequence, it would be difficult to accept, since it is not clear that there is good reason to think that conscious perception encompasses knowing how to access missing detail but does not encompass the detail itself. Moreover, if the theory deflates qualia so much that phenomenal consciousness is no longer supposed to exist at all, the sensorimotor account loses much of its explanatory advantage when compared with accounts of neural-correlates of consciousness, which themselves invite the suggestion that there is nothing more to phenomenal consciousness than being in a particular neural state. Although it is a conceptual possibility that we are radically deceived about the existence of perceptual consciousness, the claim is not a good fit with the sensorimotor theory.

There are, however, more promising ways to make sense of the claim that everything is hidden. At times, in particular in his description of P-property experience, Noë (2004) implies that understanding which movements will bring an object (or part of an object) into view is more than merely knowing how to access missing detail. It is equivalent to knowing what that object is actually like. As we saw earlier, knowing what movements to make to face an object involves knowing the object's relative location, and having the ability to reach out and grasp it, for instance, implies knowing something about its perspectival shape. With this in mind, presence-as-access could entail, in a disanalogy with the online newspaper case, that perceivers have somewhat detailed prior expectations about what the results of their bodily engagements with the environment will be. These expectations may be enough to explain the presence-in-absence (rather than mere absence) of determinate features of the visual environment. Moreover, there could be nothing more to perceptual experience than these expectations.

The claim is plausible when we consider the analogy with touch, favoured especially by O'Regan (1992; 2011). You close your eyes and pick up an object, for instance a bottle. Your feeling of holding a whole bottle, despite only being in sensory contact with a small part of it, is explained by expectations about the results of possible movements. At the same time, you have a less specific expectation that reaching out to touch other parts of the environment will reveal the presence of an assortment of further objects. The presence of the bottle, and the more general presence of a detailed environment, are genuine examples of conscious presence and are each accounted for by sensorimotor mastery, even though they are slightly different phenomena. It might seem that there is also a narrower kind of presence which does not consist of sensorimotor expectations but merely of a conscious sensation caused by parts of the environment you are in sensory contact with, for example, the feeling of pressure on your fingertips where they are in contact with the bottle. There is certainly a difference between your awareness of the parts of the bottle that are in contact with your skin and those that are not - I will examine Noë's explanation of this in a moment. What is in question, however, is that the experience of pressure on your fingertips is anything other than the expectation that given certain movements the input to your sense organs (O&N, O'Regan, 2011) or your alignment to the object (Noë, 2004; 2011) would vary in systematic ways. The sensorimotor

account contends that once we discount these expectations, there is no raw feel - as Noë would put it, no qualitative residue - remaining.

One might be tempted to think that perceptual consciousness divides into a sensational component, which might be generated by inputs from the eyes and other sense organs, and a separate intentional component that gives sensations their world-directed content. Noë denies that sensations account for phenomenal character. He says: “We can’t explain the qualitative features of experiential states (veridical or nonveridical) by supposing that they depend on the instantiation of sensations in consciousness” (2012, p. 63), and elsewhere: “Where there is an object, but no [sensorimotor] understanding, there is nothing that even rises to the level of being misleadingly like perceptual consciousness; there is only, in effect, blindness” (2012, p. 25). This might be taken to allow that perceptual experience does factor into distinct components, namely sensation and sensorimotor mastery, but that sensorimotor mastery is a necessary condition for perceptual experience.

This would be one way of explaining why your experience of the parts of the bottle that are in contact with your skin seems different from your experience of the parts of the bottle that are not. However, given that sensations cannot explain the phenomenology or intentionality of perceptual experience without sensorimotor mastery, it is hard to see why they should play a role once sensorimotor mastery has entered the story. Thus we should take it that the feel of the parts of the bottle in contact with your fingertips is explained by your grasp of how your relation to the bottle would change in line with movement, and does not depend on anything distinct from Sensorimotor Understanding. The present story makes a good fit with the claim that Sensorimotor Understanding is realised by the deployment, subpersonally, of representations of SMCs (a claim found, for instance, in a proposal by Seth, 2014, more on which in chapter 4). It is also compatible, in principle, with the suggestion that perceptual consciousness is enabled by nothing more than a non-representational dynamical system.

There is a slightly different way to understand the claim that all presence is presence-in-absence, which suggests that conscious perception does not just depend on possessing practical knowledge about the movements you could make to bring missing detail into view, but also, in an ineliminable way, on the way that the knowledge is exercised in the

course of a perceiver's temporally-extended engagement with the environment. Noë does not unequivocally endorse this suggestion, but hints at this view by observing that perception is a temporally-extended process (2004; 2012). As Clark (2009) comments, the way the process unfolds over time may be essential to the character of perceptual awareness, meaning that there may be nothing about a perceiver's physical state considered *at a fixed moment in time* that would yield perceptual awareness with a determinate content or phenomenal character. If this were true, it would be possible to say, with reference to the perceiver's bodily dispositions, which aspects of the environment the perceiver, at a given moment, has the possibility of accessing. However, we would have to refer to the perceiver's temporally-extended engagements with the environment, which might incorporate the exercise of those dispositions, to fully account for the character of her perceptual experience. I make a suggestion along these lines in chapter 7.

I have one final point to discuss before concluding this section. Noë (2012) observes that there is a need to distinguish the kinds of perceptual presence enjoyed respectively by the front of the tomato, the back of the tomato, and objects in another room. He claims that perceptual presence admits of degrees, such that the front of the tomato has more presence than the back of the tomato, while objects in the next room are not, in a strict sense, perceptually present at all. To account for this, he appeals to quantifiable properties called movement-dependence and object-dependence, each of which are necessary for full-blown perceptual presence to occur. The former is the degree to which movements by the perceiver 'manifestly' change the perceiver's sensorimotor relation to a feature of the outside environment, while the latter is the degree to which movements by the object do the same.

Assume, for the time being, that a 'manifest' change in the sensorimotor relation is one that coincides with a change in retinal stimulation. I will return to this point in a moment. Noë observes that small movements by the eyes or the tomato are enough to manifestly change the sensorimotor relation between the perceiver and the front of the tomato, meaning this relation bears a high degree both of movement-dependence and object-dependence. A small rotation by the tomato is enough to manifestly change the perceiver's sensorimotor relation with the back of the tomato, but something greater than an eye movement is required by the perceiver to effect a manifest change in this case - for instance, she might have to stand

up and walk to the opposite side of the table. This means that the subject's sensorimotor relation to the back of the tomato bears a high degree of object-dependence and a more limited degree of movement-dependence. This accounts for the lesser perceptual presence of the back of the tomato. The perceiver's sensorimotor relations to objects in other rooms, or other cities (etc.), bear only negligible degrees of object-dependence and movement-dependence, since dramatic movements by the perceiver or object are required to effect changes in the retinal stimulation the subject receives from those objects. The degree of object- and movement- dependence in these cases is sufficiently low to entail that the objects are not perceptually present at all, although they may, as we have seen, bear a kind of quasi-perceptual presence.

It is apparent that Noë has to stipulate that the movements in question cause 'manifest' changes in the character of the sensorimotor relation, since a perceiver stands in a sensorimotor relation to every object in the world, each of which is equally prone to changing as the object or perceiver moves. However, he does not explain what it takes for a change in the relation to be manifest in the appropriate sense. Notice that in some cases, we might understand the way movements cause sensorimotor relations to change even when they are relations with unperceived objects. For example, you could be in Europe walking due north and reflecting on the fact that every step forward takes you an equal distance closer to the North Pole. In the event that the position of the North Pole suddenly changes, a compass needle will be enough to show you how your sensorimotor relation with the North Pole has changed. This means your relation with the North Pole satisfies the criteria for object-dependence and movement-dependence, somehow construed. This would not, of course, count as a case in which the North Pole is perceptually present. This is not intended as a counterexample to Noë's claims, but simply to underline that an explanation must be given of what it takes for a change in a sensorimotor relation to be manifest in a distinctively perceptual sense.

This is important, because the invocation of a perceptually manifest change to the sensorimotor relation invites a misunderstanding of the sensorimotor theory. Specifically, it invites the faulty view that the changes to the sensorimotor relation become manifest first, and that this subsequently causes the perceiver to revise their understanding of the

sensorimotor relation. One understanding might be that objects cause us to have conscious sensations, and that the sensations come to be intentionally-directed when we understand how they change in line with movement. As I have noted, Noë rejects this understanding of the sensorimotor account: rightly so, as it would undermine the insight that phenomenal character can be accounted for naturalistically by nothing more than a skill-driven sensorimotor relation to the outside world. An alternative understanding of the sensorimotor theory might say that the world is, in fact, represented pictorially in perceptual consciousness, and that changes to this pictorial representation enable you to revise your understanding of the sensorimotor relations you stand in. This is no good, because sensorimotor understanding is required for perception to be intentionally-directed in the first place, and we could not therefore appeal to the world-directed content of pictorial representations to account for sensorimotor understanding.

The proper understanding of ‘manifest’, I suggest, instead specifies that the perceiver does not consciously access whatever change in input causes the change in the sensorimotor relation to be manifest. On this basis, we should take it that the relevant input is a subdoxastic sense input. The upshot, given that movement-dependence and object-dependence are necessary conditions for the perceptual presence of P-properties and objective properties, is that a necessary condition for perceptual presence, at the personal level, is that subdoxastic sense inputs originating from the objects and properties consciously perceived must currently be prone to exhibiting patterns of movement-related change.

Notice that the claim that perceptual experience does not factor into an occurrent and a merely potential part is, on the present understanding, compatible with the outside memory thesis (O’Regan, 1992), which claims that visual processing involves an occurrent part, namely the retinal image, and a potential part, namely expectations - represented, or physically-constituted in a non-representational way - about the consequences of possible movements. The current personal level account can be taken to imply that this distinction between the occurrent and the merely potential does not manifest itself in a straightforward way to the conscious agent. This makes sense when we consider that appeals to the character of retinal stimulation are not sufficient to explain the phenomenal character or intentionality of perceptual consciousness. In light of this limitation, Sensorimotor Understanding is all

there is to go on. Noë's notion of presence-as-access is in principle compatible, unlike the outside memory thesis, with the existence subpersonally of filling-in and pictorial representation, since it does not need to claim that there is a one-to-one mapping between perceptual experience and the physical processes that enable it. Indeed, Pessoa et al. (1998) and Noë and Thompson (2004) actively deny that there is such a mapping between subpersonal content and the contents of visual experience. Notice, all the same, that presence-as-access does not do anything to invite the view that filling-in or pictorial representation take place subpersonally, and makes a very good fit with the outside memory thesis.

3.5 - Perception and Dynamical Entanglement

According to one current within sensorimotor theory, associated in particular with Hurley (1998), perception should not be understood subpersonally as a matter of information processing, but as a dynamic process comprising a messy entanglement between different parts of the brain, non-neural body and outside environment. Hurley contrasts her view of perception and action with an orthodox conception she calls the 'classical sandwich', which claims that the mind decomposes into discrete subpersonal systems responsible, respectively, for perception, thought, and action. Perception and action are the respective top and bottom layers of the sandwich, and thought makes up the filling. Thus perceptually-guided action, on the classical sandwich conception, takes the form of a 'sense-represent-plan-move' cycle (Wheeler, 2005). The visual system uses retinal inputs to construct a representation of the outside world, in a manner broadly in line with the approach offered by Marr (1982). This representation serves as an input to an informationally encapsulated and functionally distinct central cognitive system. The central system chooses a course of action and sends instructions to a separate motor action system, which is responsible for generating the appropriate motor movements. Fodor's (1983) account is a good example of this kind of framework.

One's preferred conception of the causal processes involved in action and perception can be laid out diagrammatically with sense inputs at the top and motor outputs at the bottom. Hurley suggests that the classical conception endorses 'vertical' modularity, since the

modular layers responsible respectively for perception, cognition and action are stacked vertically on top of one another, with perception at the sense input end and action at the motor output end. Hurley agrees that the mind decomposes into a number of informationally encapsulated systems geared toward differing tasks, for example different kinds of perceptually-guided action. The layers in her conception, however, are stacked side-by-side, meaning the boundaries between sense inputs, motor outputs and central cognition are non-existent or at least considerably more porous than conceived by the classical sandwich view. Thus one distinctive feature of perception in Hurley's account is that it depends constitutively on fluidly integrated causal processes running between sense inputs and motor outputs.¹¹

Hurley suggests that the causal flow runs from sense input to motor output and, importantly, loops back to sense input again, as body movements alter the character of the sensory stimulation received from the environment. In addition, she claims that the causal flow runs in the opposite direction too, as efferent copy from the motor outputs makes a causal impact upstream. She states that these points distinguish her account from behaviourism, which conceives of cognition as depending solely on unidirectional relations between sense inputs and motor outputs. Thus Hurley claims perception and action depend constitutively on messy dynamical processes, carried out without internal representations¹², and involving a complex entanglement between different parts of the brain, extra-neural body and environment. In addition to claiming that the substrate of perceptual consciousness is wider than sometimes thought, incorporating motor regions of the brain and perhaps the non-neural body and environment, Hurley claims that perceptual experience depends constitutively on the way the interactive process unfolds over time. Thus she rejects temporal

¹¹ Chemero and Cordeiro (2000) object to Hurley's continued use of the words 'input' and 'output' on the ground that on the dynamical conception, which they and Hurley endorse, there is no principled reason for conceiving of them as inputs and outputs. Still, we can describe them as inputs and outputs for heuristic purposes while acknowledging that this may not reflect a genuine functional distinction.

¹² As with 'inputs' and 'outputs', Hurley retains an apparently orthodox commitment by endorsing the vehicle/content distinction, implying that the entire processing loop could be considered one big representation. This is nonetheless quite different from the wholly internal representations endorsed by the classical sandwich conception.

atomism, the claim that the processing can for explanatory purposes be broken down into temporally discrete chunks.

With the points in mind, perception (and action) are held by Hurley's account to be best interpreted as dynamical systems. As such, they have the potential to be modelled using dynamical systems theory, which uses difference or differential equations to describe the ways a mechanistic system evolves over time. Hurley's account is just one example of a broader 'dynamicist' approach to cognitive science, associated in particular with van Gelder (Port and van Gelder, 1995; van Gelder, 1998) and Thelen and Smith (1994) and encompassing work in biological enactivism (e.g., Thompson and Varela, 2001) and ecological perception (Chemero, 2009). All endorse the suggestion that cognitive activities can be best understood, and in particular formally modelled, as temporally-evolving, dynamical processes, usually without reference to representation.

Aside from Hurley's account, canonical work on the sensorimotor theory is by and large non-committal about the dynamical approach. O'Regan (2011) makes no mention of dynamical systems, appearing, in spite of his account's emphasis on interactive engagements with the outside environment, to prefer an information processing stance. Noë (2004) only claims that perception *might* depend constitutively on dynamic input-output-input loops. The dynamical view is given a fuller endorsement, however, by other work on the sensorimotor theory, including Hurley and Noë (2003), Gangopadhyay and Kiverstein (2009) and Burhmann, Di Paolo and Barandiaran (2013).

There are numerous available motivations for endorsing the claim that perception depends constitutively on an embodied dynamical process. One such motivation is evidence that the phenomenal character of perceptual consciousness is partly determined by motor output activity. Standard representationalist models do not expect perceptual experience to be non-instrumentally dependent on activity by the brain's motor regions. If perceptual experience depends constitutively on motor activity, it plausibly depends on entire loops of interaction between brain, body and outside environment. Hurley (1998) appeals to an experiment by Kohler (1962) featuring goggles in which the left half of each lens was tinted yellow and the right half tinted blue. After wearing the goggles for several weeks, the subjects adapted to the coloured lenses, meaning that the visual scene came to appear normal and

untinted. After removing the glasses, the subjects reported that moving their eyes to the left caused objects to appear with a yellow tint, while moving their eyes to the right produced a blue tint. Hurley claims that this supports the view that motor output states partly drive visual experience and determine its phenomenal character.

The dynamical approach can also be motivated by the idea of the world as an outside memory. Although O'Regan does not pursue the outside memory idea to this conclusion, the idea that perception is best characterised as a dynamical system is a natural extension of the claim that perception depends on embodied interactions and not solely or primarily on internal information processing. Thus the arguments provided by O'Regan and others in favour of the outside memory view, for example the empirical work on change blindness and inattention blindness, give some evidential support for a dynamical systems version of the sensorimotor theory.

Elsewhere, the dynamicist approach is strongly associated with accounts that claim internal representations do not exist at all or cannot do the explanatory work required to bridge the explanatory gap between consciousness and the physical. As Chemero (2009) comments, dynamical systems models can frequently be given a representationalist gloss, but such a gloss is not necessarily required for the models to do the explanatory work they do. This contrasts with orthodox information processing accounts, which typically make representation an essential posit. Hutto and Myin's account (2012), which claims on conceptual grounds that subpersonal representational content probably does not exist, thus provides support for the dynamical systems perspective. Biological enactivists (Varela, Thompson and Rosch, 1991; Thompson, 2007) have a further reason to endorse dynamical systems, namely the view that cognition is continuous with the self-organising dynamical systems they hold are constitutive of life.

Chemero (2009) claims that the merits of a nonrepresentational dynamical systems approach must be judged by its explanatory success in making sense of empirical data rather than by metaphysical considerations about content of the kind endorsed by Hutto and Myin. However, Silberstein and Chemero (2012) offer a metaphysical-motivation of sorts for adopting a dynamical systems perspective, claiming that it helps motivate a direct realist account of phenomenal qualities, which in turn means it promises to do a better job than

competing accounts of explaining phenomenal qualities physicalistically. I will return to this point in the next section.

3.6 - The Extensive Conscious Mind

The sensorimotor account, in some variants, advocates one of the more extreme departures offered recently from the mainstream view that the mind is located purely 'in the head' and does not need to be explained by appeal to the body or outside environment. The orthodox position, to which the sensorimotor theory sets itself in opposition, is sometimes traced back to Descartes, who claimed that the mind was an immaterial entity distinct from the body, with which it interacted via a prescribed channel at the pineal gland. Since the mind, on the Cartesian view, is not even physical, it is in this sense more disengaged from the body and physical environment than it ever is in contemporary cognitive science. However, orthodox physicalist accounts of mind continue to subscribe to what Wheeler (2005) describes as Cartesian psychology, an approach which implies, amongst other things, an analogous separation between the neurally-realised mind and the extra-neural body. Descartes was, in fairness, careful to claim that the mind is intimately joined with the body and not detached, as he put it, in the way a pilot is from their ship (1647/1985b). Wheeler, attempting to do justice to the position, suggests this claim can be reconciled with Descartes' dualism if we regard it as a principle of Cartesian psychology that the mind makes use of informational content specifying the body's state, but that the way it uses this information to generate intelligent action can be fully characterised without reference to the body. This principle, drawn from a nuanced version of Cartesian psychology, is explicitly endorsed by a number of recent, largely orthodox accounts of cognition, which hold that the mind supervenes purely on the brain but frequently trades in representations of bodily states (e.g., Goldman and de Vignemont, 2009; Rupert, 2011).

Drawing a less kind parallel between Descartes and contemporary accounts, Dennett (1991) issued a well-known warning against being tacitly committed to Cartesian materialism, the view that the mind is unitary and its physical substrate confined to a circumscribed area of the brain - an outlook that must be avoided on pain of committing the homunculus fallacy.

Present day theories often maintain an orthodox position while appearing to avoid this mistake. Rupert (2011) is again a good example, since it is the central thesis of his ‘massively representational mind’ that representations are manifold throughout the brain, although he maintains that cognition is overwhelmingly brain-bound. Noë (2004) argues that brain-bound accounts cannot avoid the Cartesian mistake altogether, since they continue to endorse a conception of the brain itself as a homunculus subsisting within the wider confines of the person as a whole. This need not be viewed as a mistake, so long as you are happy to endorse the view that in an important sense, you are your brain. This is nonetheless strikingly similar to Descartes’ claim, with reference to his immaterial mind, that “I am in the strict sense only a thing that thinks, that is, I am a mind or intelligence or intellect or reason” (1647/1985b, p. 18). The moral is that although orthodox theories are often deliberate in avoiding the extremes of Cartesianism in a crude rendition, they retain Cartesian commitments.

One point of departure from the Cartesian orthodoxy occurs with the thesis known as vehicle externalism. To explain the generic position, it will help to compare it with content externalism (Putnam, 1975; Burge, 1979). This longer-established position holds that mental content is not fixed solely by things in the head. In Putnam’s classic thought experiment, Oscar, who lives on Earth, has an exact physical duplicate, Twin Oscar - who lives on Twin Earth. Twin Earth is identical to Earth in every respect, except the clear liquid which fills rivers, lakes, and has all the other surface features we associate with water, has the chemical composition XYZ rather than H²O as it does on Earth. Oscar and Twin Oscar live in medieval times and are ignorant of the microstructural properties of water. But when Oscar and Twin Oscar each utter the word ‘water’, the differing properties of the natural kinds they are in causal contact with entails that they express different thoughts.¹³ The moral is that identical vehicles can bear non-identical contents, because content is fixed in part by properties of the outside environment. Vehicle externalism bears comparison with content externalism, although it is a quite different position. Importantly, content externalism is not

¹³ When Putnam (1975) formulated the view, he intended it to apply to linguistic rather than mental content, but the thesis came later to be widely applied by Burge (1982) and others to the thoughts underlying the utterance and not just the words themselves

generally thought to threaten the view that the vehicle of thought and experience is wholly neural. Vehicle externalism, by contrast, contends that the physical substrate of contentful mental states is sometimes outside the head.

Vehicle externalism, generically, decomposes into a number of quite distinct and, indeed, frequently incompatible strands. One such position is the extended mind thesis (Clark and Chalmers, 1998; Clark, 2008). Proponents of the extended mind frequently endorse functionalism (Clark, 2008; Wheeler, 2010), and hold that although cognition frequently involves the construction and deployment of representations, typically in the brain, it sometimes makes use of representational vehicles outside the skin. The most famous example, from Clark and Chalmers (1998), uses a 'parity'-based heuristic to lend support to the extended mind. Two protagonists, Otto and Inga, go about their daily business. But Otto, unlike Inga, suffers from mild dementia, meaning that unaided he suffers from severely impaired short term memory. To compensate, he makes use of a notebook which he keeps with him at all times, frequently noting down relevant facts and occurrences from his daily life, and calling upon the notes automatically and non-critically as needed. By interacting with the pen and notebook, Otto is able in a relatively smooth way to replicate the mental function for which Inga relies on her neurally-encoded short term memories. Clark and Chalmers suggest that because representations in the notebook, in Otto's case, function in an appropriately similar way to the relevant neural representations for Inga - hence the parity - the pen and notebook can be regarded as forming part of an extended cognitive system.

The extended mind thesis is one way of breaking from the tradition which says that the physical substrate of cognition is purely internal. However, the thesis is conservative in several key respects. It maintains a commitment to internal representation, and indeed uses the idea that cognition frequently depends on internal representation to motivate the parity-based claim that Otto's cognition extends, given that he uses representations in the notebook in the way other cognisers use internal representations. Moreover, it is committed only to the claim that unconscious mental processes sometimes extend beyond the skin, and proponents of the extended mind are sometimes explicit in denying that the substrate of consciousness extends (e.g., Clark, 2009).

The extended mind contrasts with a quite different vehicle externalist approach

associated with sensorimotor theory and other dynamical-systems oriented accounts of perception, which claim that the substrate of perceptual consciousness incorporates the extra-neural body and environment in addition to the brain (Hurley, 1998; Noë, 2004; Rowlands, 2010; Silberstein and Chemero, 2012). This approach, which I focus on for the remainder of the section, is sometimes called the ‘extended conscious mind’, a label intended to capture its resemblance to the extended mind. In fact, this label may not have the appropriate resonance, since it suggests a continuity with the extended mind thesis that does not reflect its quite different motivation. Vehicle externalism, strictly speaking, is a problematic label too, since some of the views falling under the present approach reject the claim that cognition, and specifically perception, involves content at all, and hence that there is a vehicle/content distinction to be made in this case. After Rowlands (2010), we could recast the distinction between vehicles and contents as a distinction between vehicles of consciousness and the conscious states that supervene on them.

A better label is perhaps suggested by Hutto and Myin (2012), who describe their approach to cognition as the *extensive mind* to underline their view that because there is no content there is, in the first place, no theoretically relevant inner-outer boundary between the brain and its surroundings from which cognition might extend. Hutto and Myin are sympathetic but not fully signed-up to the claim that perceptual consciousness supervenes on extra-neural activity. To make sure it has all the correct resonances, let us call the present view the *extensive conscious mind* (ECM). The core principle underlying ECM is the claim that perception is best explained by appeal to dynamical systems rather than computation. Although this claim is motivated in varying ways, and in turn supposed to justify ECM on the basis of varying further premises, adherents of ECM consistently endorse the principle that perception is best understood as a dynamical system, in principle describable using differential equations, incorporating features of the brain, body and outside environment.

A large part of ECM’s appeal is its distinctive way of accounting naturalistically for consciousness and phenomenal qualities. Because the objects of perception are themselves constituents of perception, phenomenal qualities can be accounted for non-mysteriously by appeal to physical properties the objects bear in relation to the perceiver’s body - for example, SMCs. One place ECM is endorsed is in Silberstein and Chemero (2012, henceforth S&C),

who advocate a version of the ecological approach to perception. The approach differs from sensorimotor theory by claiming that the features we are responsive to in perception are not SMCs but affordances. Like SMCs, affordances are relations between a perceiver and their environment, and can be straightforwardly specified from a third-person perspective. Unlike SMCs, they describe possibilities for bodily movement rather than sensory consequences of possible movement. Notwithstanding this difference, S&C's arguments in favour of ECM are compatible with the sensorimotor theory.

S&C emphasise that ECM can dissolve the hard problem of consciousness, by which they mean the problem of accounting naturalistically for specific phenomenal qualities *and* the more general question of why there is consciousness at all. They take inspiration from a point made by Ryle (1979), who identified a dilemma faced by theories of discursive thought. Ryle claimed the theories were prone to pivot between two unsatisfactory positions, either casting thought as 'nothing but' inner speech, or casting it as 'something else as well'. He illustrated the problem by analogy with a coin. It would be wrong to claim that a coin is 'nothing but' a metal disc, but wrong to claim it consists of a metal disc plus some other entity as well. S&C argue that the problem facing theories of consciousness takes a similar shape. Proponents of orthodox approaches to cognition sometimes claim that phenomenal consciousness is 'nothing but' a certain kind of neural activity - an unsatisfying position which also comes under target in O&N, as I discussed in section 3.2.2. Alternatively, they accede to a dualist impulse and claim, like Chalmers (1996), that consciousness is a non-physical 'something else as well'.

S&C blame computationalists for the apparent dilemma, offering the following diagnosis: Computationalism, which as S&C take it entails the representational theory of mind, states that computation is sufficient for cognition. Computationalists remain loath, however, to accept that metal and plastic artifacts could be conscious merely by virtue of performing appropriate computations. They are therefore forced to separate consciousness from cognition, and for this reason account for consciousness separately, either deflating it to brain state or inflating it to something non-physical. To avoid the dilemma, S&C reject the representational theory of mind and instead propose that the properties to which we are intentionally-directed in perceptual consciousness are themselves constituents of perception.

The fact that the properties we perceive, affordances, are specified with egocentric reference to the perceiver is supposed to account for perception's intrinsic subjectivity, and therefore consciousness in general.

Although S&C's solution is appealing, it is worth taking note of a limitation they share with O&N. It is true that the representational theory of mind is not necessarily committed to the claim that consciousness, like cognition, is best explained by appeal to representation. However, there are also representationalist proposals which take seriously the claim that contentful states, meeting certain other functional criteria, are sufficient to account for phenomenal consciousness (e.g., Lycan, 1987; Dretske, 1995; Tye, 2000). Accounts in this category accord with S&C's demand that consciousness is neither 'nothing more' than brain state nor 'something else as well', and like S&C they attempt to avoid the dilemma by identifying consciousness with cognition. However, they do not use ECM to make this work. ECM has one key advantage compared to representationalism about phenomenal qualities, namely that it avoids much of the latter's explanatory burden. ECM does not have to offer (or hold out for) a naturalistic account of content, a project which faces serious conceptual obstacles (Hutto and Myin, 2012; see section 6.4). Notice that reducing phenomenal character to content only helps naturalise phenomenal character if the content can itself be naturalised. Moreover, the success of representationalist accounts of phenomenal character is in part dependent on the sort of content involved. For instance, externalistically individuated content faces a distinctive challenge in light of Twin Earth-style cases, as we will see shortly.

Hurley (1998), offering a different argument for ECM, endorses the vehicle/content distinction, arguing that the vehicle of conscious content extends beyond the brain. She claims that vehicles of content play an essential theoretical role in explaining the relation between the mental and the physical. Drawing a distinction between type-explanations and token-explanations, she claims the former explain why there is mentality at all. For example, in line with Millikan's (1984) biosemantics, facts about natural selection and the ancestral environment might explain how certain types of physical state bear content. To explain a subject's present mental state, we must also appeal to presently occurring states and processes. We are now offering a token explanation, because we are specifying which physical states and processes token the relevant type. Thus Hurley claims that when we talk about a vehicle of

content, we are attempting to offer a token-explanation.

If one wanted to deny that perceptual experience is contentful or depends on content-bearing vehicles, the vehicle/content distinction could be recast as a relatively harmless vehicle/experience distinction. This avoids the need to stipulate that experience involves truth or correctness conditions. However, maintaining the reformulated distinction is only valuable if phenomenal states are multiply realisable, meaning that non-identical physical processes can yield phenomenologically indistinguishable experiences. S&C maintain that the dynamical systems which constitute perception are ‘soft assembled’ from varying physical components, sometimes wholly internal and at other times incorporating varying parts of the outside environment. This implies that perception, in general, is in one sense multiply realisable. However, the point here is that the vehicle of perceptual experience in general can vary, bringing with it attendant differences in perceptual phenomenology, and not that an identical experience could supervene on a range of different vehicles. The sensorimotor theory, on the gloss I argue for later in the thesis, allows that perceptual experiences may in a weak sense be multiply realisable, given that phenomenal character could stay the same across perceivers with limited differences in body morphology. However, the multiple realisability that applies in this case is too constrained to make a vehicle/experience distinction particularly useful, especially when contrasted with the functionalist thesis that phenomenology could stay the same despite radical differences in embodiment.

Hurley’s arguments in favour of vehicle externalism are equally applicable to other versions of ECM. Conceding that neural processes are likely to be sufficient for consciousness, she challenges the assumption that suitable attendant outer states are not themselves necessary for the neural processes in question to occur. The ‘duplication assumption’, as she calls it, is a standard background assumption underlying the twin test for content externalism. Recall the Twin Earth thought experiment from Putnam (1975). Block (1990b) offers a similar thought experiment intended to test representationalist accounts of phenomenal character (see chapter 5). Suppose there is a planet called Inverted Earth, which is like Earth except the colours of objects are all inverted. A subject travels from Earth to Twin Earth and puts on colour-inverting goggles. As a result, the subject’s phenomenal states remain as they were on Earth, even though, according to content externalist accounts, the content of the subject’s

perceptual states must differ. According to Block, this demonstrates that content externalist variants of representationalism about phenomenal qualities fail. His own view is that we should appeal to some non-representational property of the brain to explain phenomenal character. Hurley's argument undercuts this debate by casting doubt on the claim that twin tests are possible. If successful, it challenges Block's view that we must only appeal to the brain, since it sheds doubt on the claim that phenomenal states can stay the same while the environment radically changes. It challenges the content externalist proposal, since it denies that the vehicle can stay the same while the content changes. Notice that it thereby offers a response to representationalism about phenomenal qualities absent from S&C's more recent piece.

Hurley observed that the duplication assumption becomes problematic when we apply it to realistic rather than toy cases. In *Consciousness in Action*, she offers a series of Twin Earth-style scenarios which challenge the assumption in progressively demanding ways. A first concerns a case of colour inversion in which all objects that are red on Earth are green on Twin Earth but *not* vice versa. Because the transformation between Earth and Twin Earth is not information preserving, the subject on Twin Earth is unable to make the same range of discriminations as the subject on Earth, therefore their neural states are not duplicated and the duplication assumption fails. This shows that even for cases where the differences between the Earths are limited to colour and not shape, the duplication assumption only holds for a heavily restricted set of cases.

A further scenario features a planet called Mirror Earth in which everything is reversed from left to right, so that everything is flipped on a vertical axis. The Mirror Earth subject wears left-right inverting goggles, meaning their retinal input, right now, is identical to the state it would be in on Earth. However, when the twin moves their right hand, the visual feedback they receive duplicates the visual feedback they would receive from moving their left hand, not their right, on Earth, and therefore duplication fails. To correct for this, we could imagine that a device is attached to the twin's peripheral nervous system, intervening in their motor outputs and proprioceptive inputs, such that the same instruction causes the subject to move their left hand, and hence receive the necessary visual feedback, while receiving the proprioceptive input they previously received after moving their right hand. This saves the

duplication assumption, for the time being, although it depends on confining duplication to the central rather than peripheral nervous system.

Lastly, there is El Greco Earth, in which everything is vertically elongated compared to Earth, such that the counterparts of spheres are egg-shaped. This leads to a challenging asymmetry between the respective Earths. When an Earth subject nudges a sphere, its orientation stays the same - perhaps it rolls a little to the left or right. The same action on El Greco Earth causes the egg-shaped object to fall over. Because of the asymmetry between the effects of motor movements on visual input in subjects on the two worlds, duplicating the Earth subject's internal states on El Greco Earth would require a vastly complicated device to correct the visual inputs in different situations. On a similar theme, Hurley remarks that real life Virtual Reality simulations depend on heavily restricting the range of possible movements afforded to the perceiver. The more freedom of movement the subject has, the greater their ability to disambiguate visual inputs by moving around, which means, claimed Hurley, that creating a fully realistic virtual reality simulation pushes toward the limit of what is, even in principle, physically possible. Noë (2004) makes a similar point, asserting that a completely realistic virtual reality simulation cannot be much less complex than the world itself.

The upshot, for Hurley, is that once it is established that the duplicationist must go to byzantine lengths, at the bounds of physical possibility, to preserve duplication with regard to a restricted area of the brain, the vehicle internalist thesis comes under threat. For one thing, there is no guarantee that the same area of the brain will be duplicated in each Twin Earth scenario. For another, even if some part of the brain can be consistently duplicated, without some strong motivation, the insistence that this region of the brain alone is the vehicle of consciousness appears arbitrary rather than principled. An important feature of this discussion is that it sets out to challenge duplication for a range of realistic cases. This does not mean denying that duplication is possible, even straightforward, in *some* cases. Establishing it in some cases is insufficient to show that perception in general supervenes only on neural activity. Hurley observes that in cases that are restricted enough to make duplication straightforward, for example because we imagine the subject's sensorimotor capacities are highly constrained, it is doubtful whether the subject enjoys perceptual experience in the first place.

She also pre-empts a number of more sophisticated objections that might be offered by a vehicle internalist. One such objection is that the laws of physics might vary between Earth and Twin Earth, such that complicated devices for altering the nervous system are not required to achieve duplication. Hurley counters that the twins cannot be considered duplicates if the laws of physics vary between the cases. Another objection is that a mad scientist could duplicate a subject's neural states in a different environment by intervening to alter them frame-by-frame. Hurley rejects this, because she rejects temporal atomism, the view that perception depends on a succession of moment by moment states rather than a smoothly continuous process.

The success of Hurley's argument, it should be noted, depends in part on endorsing a dynamical systems approach to perception. Indeed, Clark (2009) observes that it depends on more than just the claim that perception involves a temporally-extended, dynamical process. Such a process may merely be instrumental in causing the activation of a brain state which itself generates a perceptual experience in an instant, without requiring temporal extension. Alternatively, the brain state may need to persist through time, but it may not itself need to dynamically evolve through time in a way that would make the contribution of extra-neural states essential. Thus Clark suggests that Hurley's view must endorse a conception he calls Deep Dynamic Entanglement + Unique Temporal Signature, or 'DEUTS' (p. 980), which holds not only that perceptual consciousness is enabled by a temporally-extended, dynamic process, but also that the conscious experience is constitutively dependent on the precise way the neural as well as extra-neural states involved evolve over time.

If there are satisfactory independent grounds for thinking that perception depends on a specific functionally distinct, informationally encapsulated area of the brain, then it is no longer a prejudice to think that the same area of the brain can be duplicated in all the Twin Earth scenarios (notwithstanding the technical challenge) and that the area of the brain so duplicated is genuinely the area responsible for visual awareness. Of course, were the commitment to this view of perception motivated by a prejudice (recall the discussions of Cartesian materialism and analytic isomorphism) it would not be a suitable independent ground. Importantly, although a well-motivated version of vehicle internalism might support the duplication assumption, the duplicationist hypothesis cannot be used to motivate vehicle

internalism, as this would be getting things the wrong way round. Conversely, Hurley's argument shows that if perception depends constitutively on a dynamic, interactive process with a unique temporal signature, there is also a strong motivation for endorsing ECM.

Clark argues that if deep dynamic entanglement and a unique temporal signature are indeed essential to consciousness, this threatens a stalemate between ECM and its opponents. He argues that DEUTS cannot establish ECM, but it can expose the rejection of ECM as a prejudice. To show that ECM fails, the opponent has to identify a special property that the neural component of the processing possesses and that the extra-neural activity lacks, and show that this difference is a principled rather than arbitrary ground for thinking that perceptual consciousness supervenes solely on the neural activity. Clark claims that such a difference can be established in light of DEUTS' own appeal to a unique temporal signature. One aspect of the temporal signature that consciousness is likely to depend on is informational processing at very high speed. This, in turn, depends on the high-bandwidth available in the neural portion of the wider interaction between brain, body and environment. The interface between the brain and extra-neural body has a much lower bandwidth, suggesting it cannot do the work required to generate consciousness, and therefore serves as mere causal scaffolding (Clark, 2009). Importantly, Clark does not reject Hurley's insight that the neural processes involved in a particular conscious experience are likely to require that the subject is embedded in a specific type of environment, and hence her rejection of the possibility of duplication as illustrated by the Twin Earth-style cases.

I am not going to consider any further whether ECM finally goes through, as ECM is a mere nicety in the context of the sensorimotor theory as a whole. Equally important are the insights that lead up to ECM, including, on a dynamical version of the sensorimotor account, the claim that architecture of perception involves a dynamic, interactive process with a unique temporal signature, and the fact that the neural duplication is likely to fail. Silberstein and Chemero associate ECM's claim that the objects of perceptual awareness are constitutive features of perceptual processing with direct realism. However, it is not obvious that ECM is needed for direct realism to go through. Direct realism, including the relational realism endorsed by Noë (2012), could be motivated by the fact that perception is an interactive process in which duplication scenarios are impossible, without needing ECM too. Thus the

arguments that lead up to ECM are enough to get a dynamical version of the sensorimotor theory what it really needs.

4 - Mastering Sensorimotor Understanding

The sensorimotor theory's most crucial tenet is also the one least well-defined by the theory's flagship statements. Expanding on the work done in the previous chapter, I will in this chapter address in further detail the claim that perception depends on sensorimotor 'mastery'. Specifically, I will focus on the variety of sensorimotor mastery found in Noë's (2004) account, which I have officially labelled Sensorimotor Understanding, reflecting the terminology sometimes used by Noë. Many of the important things Noë says about Sensorimotor Understanding concern the 'presence' that arises when we exercise it (see section 3.4), and the nature of the properties it puts us in touch with, P-properties and objective properties (see section 3.2.5). In this chapter, I focus on Sensorimotor Understanding itself, with a view to addressing a lack of clarity about what it is, and how it fits into a scientific account of perception and perceptual consciousness.

Although the faculty featured in O'Regan (2011), which I label Sensorimotor Tracking, is characterised in somewhat different terms to Sensorimotor Understanding, the respective varieties of sensorimotor mastery are each loosely enough defined in the original work to leave open the possibility that they are one and the same faculty. Although my discussion will focus on Sensorimotor Understanding, the characterisation I end up with could also, therefore, describe Sensorimotor Tracking. I will continue this chapter by identifying what I regard as a key limitation with the concept presented by Noë, namely a lack of clarity about how to account for Sensorimotor Understanding physicalistically, particularly in the domain of cognitive science. In section 4.2, I address a specific ambiguity concerning the explanandum, and suggest two alternative ways of resolving it. In sections 4.3 and 4.4, I return to the scientific tractability issue by explaining the impact that the respective views of Sensorimotor Understanding should have on our understanding of sensorimotor theory as an account in cognitive science.

4.1 - 'Action in Perception' in Action in Perception

One of the most important differences between the early sensorimotor account (O&N) and

Noë's (2004) version is the leap from what we could think of as a 'bottom up' type of explanation to a 'top down' one. By bottom-up, I mean an explanation that starts by describing low-level physical states and processes that could be possessed by an unconscious machine, and expands to encompass consciousness and the personal level. By top-down, I mean an account that takes consciousness as its starting point and saves the question of how consciousness arises until later.

It will be useful if I explain Noë's motivation for this shift in emphasis. In the (2004) book, he notes that the early account focuses on how sense input signals, rather than conscious sensations, change in line with movement, observing that it therefore has no need to pre-suppose consciousness and is well-positioned to give a physicalistic account of how perceptual consciousness arises. The drawback with this strategy, Noë says, is that it comes at a cost to what he calls "phenomenological aptness" (Noë, 2004, p. 228). Expanding on this point, we might say something like this. The sense inputs featured in O&N are subdoxastic, i.e. inaccessible to the conscious agent. Even when the agent's mastery of the SMCs associated with those sense inputs is integrated with the appropriate capacities for thought and planning, the subject does not become aware of the sense inputs themselves, but of the relevant features of the outside environment. Since sensorimotor mastery, in the early account, is mastery of how sense inputs change with movement, there remains a question of which features of the outside world we consciously perceive, and how they show up for us.

The O&N account does answer these questions in part. The environmental features we perceive are those that cause the sense inputs whose patterns of movement-related change we, or our bodies and brains, have mastery over. The account also makes some observations about the phenomenal character of perceptual experience, and attempts in particular to resolve the problem of supposedly irreducible qualia by suggesting that there is nothing more to experiencing the redness of red, for example, than exercising one's mastery of the sensorimotor laws associated with redness for the purpose of thought and planning. However, this solution is in danger of failing to unseat the intuitions of someone who thinks that qualia are irreducible, because it invites them to ask why exercising mastery of a particular set of SMCs has one phenomenal feel and not another - this is another of the unresolved 'why' questions about the relation between consciousness and the physical that O&N make it their

project to avoid.

The problem arises because the O&N version of the sensorimotor theory cannot convincingly claim that perception *feels* like exercising sensorimotor mastery, given that the sense inputs which are the targets of the mastery remain subdoxastic. If perception does not feel like the exercise of sensorimotor mastery, the endorser of irreducible qualia will continue to hold that the relation between sensorimotor mastery and conscious feel is only contingent, meaning that the conscious feel is therefore nonidentical to the associated sensorimotor mastery. The sensorimotor theorist could respond that perceiving does feel like exercising sensorimotor mastery - necessarily, because perceiving just *is* exercising sensorimotor mastery - even though perceivers typically lack the ability to recognise that this is what perception feels like. However, if the sensorimotor theory concedes that conscious perception does not feel *recognisably* like the exercise of sensorimotor mastery from the subjective point of view of the perceiver, it finds itself deprived of powerful reason for thinking the sensorimotor theory is true, namely that the theory makes it impossible to conceive of a given act of perceiving feeling other than the way it feels. In any case, the early sensorimotor account simply does not go into as much detail about visual phenomenology as it could.

Like the original version of the theory, Noë (2004) claims that conscious perception is identical to the exercise of sensorimotor mastery. In a progression from the earlier account, however, he implies that introspection reveals that exercising sensorimotor mastery is recognisably what conscious perceiving feels like. This is useful because it provides a further reason for thinking that the act of consciously perceiving, with its distinctive qualitative feel, *just is* the exercise of a variety of sensorimotor mastery and nothing more. As we saw in the last chapter, the account claims that absent features like the back of a tomato are experientially present, in absence, because we know which movements would bring them into view. It also claims that objective properties, like the circularity of a plate, show up in visual experience as being mediated by the perspectival properties that presently obtain, which we implicitly understand will change in specific ways in line with possible movements.

In order to introduce these elements, Noë has to present sensorimotor mastery as a feature *of* perceptual consciousness rather than a faculty that enables perceptual consciousness. Because P-properties and objective properties are consciously accessible, it is natural to think

the faculty we use to access them is conscious, even though the understanding in question is implicit. Although Noë downplays the importance of the personal/subpersonal distinction, it is plausibly a requirement of Noë's account that the faculty we use to master P-properties and objective properties is itself a personal level faculty, since he explicitly intends the account to do justice to McDowell's (1994a) conceptualism about perceptual experience and Kant's famous maxim that "thoughts without content are empty; intuitions without concepts are blind" (CPR A51/B75), which in each case require that the *person* brings some capacities for thought to bear in experience. Noë does not say so, but it follows from this that Sensorimotor Understanding must also be amenable to a kind of self-reflexivity; not only do you consciously exercise Sensorimotor Understanding, but you have the capacity to be conscious *of* exercising it. This point notwithstanding, Sensorimotor Understanding is characterised by Noë most of the time as a conscious faculty possessed by the person rather than their brain or body, and Noë explicitly concedes that the account presupposes consciousness.

This reflects an important difference between Noë's sole-authored version of the sensorimotor theory and the others. As we have seen, the approaches to perception detailed in O&N and O'Regan (2011) focus, in particular, on processes that could be instantiated by an unconscious machine rather than a human perceiver. The accounts begin by describing a process of sensorimotor coupling sufficient for a certain kind of unconscious perceptual activity, and widen the explanation in a bottom-up fashion by introducing additional processes, the contributions of which they allege enable perceptual consciousness. Specifically, these are the integration of the lower-level process with thought and planning, and, in O'Regan's (2011) model, an additional faculty corresponding to the 'self'. Although the accounts do not give detailed proposals about how these higher faculties are physically realised, the authors intend us to suppose that the relevant faculties are things we could readily account for by appeal to appropriate physically-realizable states and processes. O&N at times suggest an information processing-style account is needed, while O'Regan (2011) explicitly endorses functionalism of a sort that means a robot could in principle possess the relevant functionally-defined faculties. Thus an important difference between these accounts and the version of the sensorimotor theory offered by Noë (2004; 2012) is that the latter pre-supposes consciousness and does not make any claim to explain it. The Sensorimotor Understanding

Noë appeals to is intended to explain something about perception and the character of perceptual consciousness, but is not thereby supposed to account for the existence of consciousness in the first place.

Neglecting to explain how consciousness in general arises is not an unreasonable limitation, as long as we think of the account as an account of perceptual consciousness rather than consciousness in general. More troubling is the fact that the account fails in great respect to commit itself to clearly defined theses about the physical processes implicated in Sensorimotor Understanding, something the sensorimotor theory could be expected to offer even in the absence of an account of how consciousness in general occurs. The omission would not be problematic if we took the account solely as an entry into philosophy of perception, which is not necessarily concerned with the physical states and processes that realise conscious perception. However, it becomes more of a problem when we consider that Noë continues to present the sensorimotor account as a rival to orthodox scientific accounts of vision. In response to Marr (1982), Noë says that: “Vision shouldn’t be thought of as a computation performed by the brain on inputs provided by the retina. What is vision? How should it be characterised computationally? This book suggests the outlines of an answer. *Vision is a mode of exploration of the environment drawing on implicit understanding of sensorimotor regularities*” (Noë, 2004, pp. 29-30). One of his points is that Marr presents the subpersonal processes that enable perception as if they were identical to the higher-level processes that perception is constituted by. This is not only meant to be a logical error, however. His objection is that the error causes Marr to get his account of the lower-level processes wrong. Indeed, Noë casts doubt on many of the principles needed to make a computational account stand up, claiming for example that perceptual consciousness depends in part of having the right kind of body (p. 25), a claim by which he intends something stronger than the mere suggestion that you need a body capable of implementing the appropriate computations.

We could simply take the sensorimotor theory as a rebuttal of orthodox scientific accounts like Marr’s without expecting it to provide an alternative account of the subpersonal processes involved in perception. It would be better, however, if the theory gave an account of the physical or physically-realised functional states implicated in Sensorimotor Understanding. Offering such an account involves answering some very abstract questions, for example

whether functionalism and representationalism are conceptually viable approaches for reducing the mental to the physical, and slightly more concrete questions, for example whether information theoretic or dynamical systems frameworks have the most theoretical utility for explaining what is going on in human brains and bodies. Even more concretely, it would be nice to know the precise mechanism involved in human Sensorimotor Understanding. For instance, if we concluded that the best way to physicalistically account for Sensorimotor Understanding is to suppose that it supervenes on the construction and deployment of neural representations, perhaps of some relatively novel kind, we would still need to know what kind of content they bear and what functional role they play.

Ward (2009), for example, has argued that the representations in question might specify possibilities for highly planned agent-level action. Elsewhere, Seth (2014) has proposed that aspects of the sensorimotor approach can be allied with predictive coding approaches to cognition, which I will explain below. In contrast with representationalist proposals like these, Hutto and Myin (2012) suggest that Sensorimotor Understanding must be explained by appeal to the agent's history of sensorimotor interactions rather than by representational content, although it remains to be properly explained how we should characterise the physical processes that underlie the exercise of sensorimotor mastery *right now* if we endorse this view. Burhmann, Di Paolo and Barandiaran (2013) have offered a way to model sensorimotor perception using a dynamic systems framework. The dynamical systems approach was also endorsed by Hurley (1998), whose account can be considered part of the sensorimotor approach, and Chemero (2009; see also, Silberstein and Chemero, 2012), who endorses a Gibsonian account of perception, but whose appeal to sensorimotor abilities offers a proposal that may usefully help do justice to the sensorimotor approach.

The founding statements of sensorimotor theory can be forgiven for not arriving with a prefabricated and fully-developed account of the physical processes involved in Sensorimotor Understanding, but if it is to convince as a serious rival to orthodox scientific accounts of vision, we need a direction of travel toward answering them. A good place to start is to have a clear working proposal about what the explanandum, Sensorimotor Understanding, actually is. In fact, even this so far is ambiguous.

4.2 - Which Direction of Fit?

In the sections that follow, I will further discuss the physical or functional processes involved in Sensorimotor Understanding, but in this section I will first address a significant ambiguity about what Sensorimotor Understanding is, namely the fact that it is sometimes pitched as Being Able to Act and sometimes as Knowing What Would Happen if You Did Act.

The former, Being Able to Act, could be understood in various ways. In a position approximating behaviourism, it could be taken to mean having a disposition or ability to act. Alternatively, in a position resembling Gibsonian ecological psychology (e.g., Chemero, 2009) or the literature on action-oriented representation (e.g., Clark, 1997; Wheeler, 2005; Ward, 2009; Ward, Roberts and Clark, 2011), it could be taken to mean knowing what actions and behaviours are open to you. Although Noë (2004) sometimes disowns these possibilities, there is much evidence that even his version of the sensorimotor theory is committed to a view of Sensorimotor Understanding as Being Able To Act. If Sensorimotor Understanding is like this, it has a *desire-like* direction of fit, because it is the state of the world, specifically the agent's behaviour (which for this purpose counts as a state of the world) that must fit with the agent's mental state, and not the other way round. Contrast this with Knowing What Would Happen if You Did Act. This calls upon (perhaps implicit) knowledge about counterfactual conditionals, and to be effective, the agent's mental state in this case must be sensitive to the state the world is in rather than vice-versa - hence it has a belief-like direction of fit.

Noë, as we will see, is sometimes explicit that Sensorimotor Understanding is not desire-like but belief-like. But to begin, let's review the textual evidence that Noë's account is committed to the idea that Sensorimotor Understanding is desire-like. Although he insistently denies charges of 'behaviourism', he endorses the desire-like direction of fit option at the very beginning of his book, where he says: "Perceptual experience acquires content thanks to our possession of bodily skills. *What we perceive* is determined by *what we do* (or what we know how to do); it is determined by what we are *ready* to do" (2004, p. 1). Throughout the book, he makes the recurrent claim that perception is constituted by the exercise of 'sensorimotor skills' (e.g., Noë, 2004, p. 90), a phrase which without some gymnastics is hard to think of as referring to anything other than capacities of one kind or

another for skillful bodily behaviour. Noë also equates sensorimotor skills with 'practical knowledge', and this, he makes a point of arguing, conforms to Ryle's (1949) view of knowing-how, which Noë suggests means it is an *ability* and does not reduce to a variety of knowing-that.

Replying to Stanley and Williamson's (2001) argument that knowing-how cannot be an ability because you can know how to do something without having the ability to do it, Noë argues that it is in keeping with ordinary language, which he takes Stanley and Williamson to be motivated by, that a ski-instructor who cannot perform the moves they teach, or a pianist who loses the use of their hands, does *not* in fact know how to perform the relevant tasks. Again, if Sensorimotor Understanding is a practical ability and not a form of knowing-that - and if it resembles the pianist and ski-instructor examples - then we must think of it as an ability to make skillful bodily movements. It is therefore desire-like; consider the incoherence of saying that Sensorimotor Understanding involves 'knowing-how what would happen if you did act'.

Noë goes on to contrast Sensorimotor Understanding - which does not necessarily draw on propositional knowledge - with linguistic ability, which does. He claims that a better analogue for Sensorimotor Understanding than linguistic ability is the ability to communicate using bodily gestures (2004, p. 90), although he does not make it explicit what one should take from this analogy. Taken in a minimalistic way, the lesson is merely that neither sensorimotor knowledge nor the ability to gesture call upon a capacity for propositionally-structured thought like linguistic ability does. Noë draws a more direct connection between Sensorimotor Understanding and gesture elsewhere in the book, however, when he suggests that sensorimotor knowledge is actually constituted, at least sometimes, by an ability to gesture. Gesture makes an appearance, in particular, in Noë's account of the skills that underlie the perceptual experience of P-properties. He claims that the ability to experience P-size - the apparent size on an object in your visual field - is equivalent to the ability, for example, to hold out an arm and represent the perspectival size using thumb and forefinger. Endorsing an explicit appeal to capacities for bodily movement made by Pettit (2003), he suggests that experiencing the motion of a ball as it flies through the air involves being able to make bodily movements that manifest a discrimination of its speed, for example by catching

the ball.

Following Evans (1982), Noë suggests that experiencing an object's location in egocentric space, for example its being 'off to the left', is identical to practically understanding the movements you would have to make to face the object head-on. Evans (1985) pointed out that it will not do to specify the relevant movements solely in terms of muscle movements, because the appropriate behaviours could include, for example walking, crawling, swimming and indefinitely many others. Evans concluded from this that we cannot derive the notion of space by specifying appropriate behavioural capacities, but must presuppose space in order to specify what the relevant behavioural capacities are, thus endorsing a notion of 'behavioural space'. In contrast, Noë endorses a notion of "egocentric, sensorimotor space" (2004, p. 89). The concepts of behavioural space and sensorimotor space appear to be the same, or closely related: both characterise the environment spatially and in egocentric reference to the subject's body. Noë does not explain what difference there is, if any between the concepts, and what he probably intends is merely to underline his difference with Evans on the subject of how perceivers master that space.

Here Noë distances himself from the idea that sensorimotor mastery is Being Able To Act. He says:

Evans (1982) sometimes seems to offer such a behaviorist account, suggesting that the experience of something as off to the left consists, as it were by definition, precisely in the possession of certain dispositions to move with respect to the thing. My claim is not behaviorist in the way Evans's appear to be. When we see a flicker on the right we know - in a practical, implicit way - that movements of the eyes to the right bring (or would bring) the flicker better into view (Noë, 2004, p. 89).

This implies that Sensorimotor Understanding is Knowing What Would Happen If You Did Act, which is to say has a belief-like direction of fit. That Noë intends Sensorimotor Understanding to have a belief-like direction of fit is underlined in an exchange between Noë and Campbell concerning the experience of P-properties. Campbell (2008) observes that Noë's account lacks an explanation of how we locate objects in sensorimotor space to begin with. Campbell's point is that Sensorimotor Understanding, taken as an implicit, belief-like mastery of how P-properties change in line with movement, is sufficient to explain the

experience of objective properties, but that we must appeal to a distinct faculty to account for the experience of P-properties. He complains that in Noë's account: "there is no attempt at all to explain what it is to perceive an aspect of an object" (p. 669). Campbell's friendly suggestion is that the sensorimotor account should explain P-property experience by appeal to neurally-realised capacities for actions such as reaching out to grasp the object in front of you. This suggestion does not appear to be very far removed from the spirit of Noë's proposal. Noë (2008) rejects this, however, on the ground that that perceptual experience is not *for* action. He responds that experiencing P-properties implies knowing what movements one *could* make to gesture to the object or to bring it into view, but emphasises that this knowledge does not require an ability or disposition to make the relevant movements.

As an aside, Noë's reply to Campbell is unconvincing. He implies that as long as we do not stipulate that P-property experience happens *prior* to objective-property experience, we need not be committed to the view that they are the result of separate faculties. But the import of Campbell's point, concerning the 'to the left' case for instance, is that in order to know which movements would cause you to face the object head-on, you need some way to recognise the state of affairs that obtains when you are facing the object head-on. You may know that *no* movements are required to face the object head-on, and that you are therefore facing the object head-on. But in this event, some account must be given to explain how you know that no movements are required. Adopting Campbell's proposal, we could propose that you know this because you have the ability to carry out a practical task like picking the object up, but Noë rejects this solution. I will examine shortly how this limitation in Noë's account might be redressed while maintaining that Sensorimotor Understanding has a belief-like and not merely a desire-like direction of fit.

Elsewhere, Noë (2004) repeatedly disavows 'behaviourism' on various grounds. He states that it is a mistake to suppose that "effects are logical constructions of their causes" (p. 118), i.e. that the behaviours that are the effects of Sensorimotor Understanding themselves constitute Sensorimotor Understanding. He claims that Sensorimotor Understanding is not a set of behavioural dispositions, but the *ground* of your dispositions (p. 88). Importantly, he notes that you do not always need to move, or even have the ability to move, in order to

perceive (p. 90).¹⁴ These are sound objections to certain positions that might be deemed behaviourist, and I will later attempt to demonstrate that a position which does claim that Sensorimotor Understanding depends on behavioural capacities can be finessed in a way that defuses these worries (sections 4.4 and 6.2).

The uncertainty around direction of fit leaves us with two opposed possibilities. On a belief-like account, Sensorimotor Understanding is a non-propositional belief-like understanding of the sensory consequences of possible movements. This possibility is difficult to square with Noë's claims that the understanding is practical knowledge or a set of sensorimotor skills, and with his claim that perceptual consciousness depends on what we are able to do. But it does justice to Noë's denial that Sensorimotor Understanding is for action, or depends on movement or the ability to move. Accounting for the physical processes that enable Sensorimotor Understanding, on this view, appears to require explaining how neural states and processes yield belief-like states, and this may require endorsing one kind of internal representation account or another, since it is not apparent what other kinds of subpersonal state could do the necessary explanatory work.

Sensorimotor Understanding, on a desire-like account, could involve knowing what movements to make to achieve various practical goals, and this knowledge might itself depend on the deployment of action-oriented representations. Alternatively, Sensorimotor Understanding may, on the desire-like account, consist of *dispositions* or *capacities* for various behaviours which are not best accounted for by reference to representation or 'knowledge' at all. Noë, at times, resists both these possibilities. In section 4.4, below, I will argue that one can avoid the most important worries associated with behaviourism while nonetheless taking behavioural capacities to be necessary and sufficient for the existence of Sensorimotor Understanding. To this effect, I will attempt to establish that Sensorimotor Understanding exists by virtue of the subject's possessing an appropriate set of behavioural capacities, but that it 'grounds' the capacities rather than being identical to them. This thesis, as will become evident, is compatible with the further claim that representations are required to explain how

¹⁴ Noë's other objection is that behaviourism denies the existence of phenomenal consciousness (2004, p. 32). However, it is not obvious that explaining experience by appeal to behavioural dispositions necessarily entails denying consciousness.

the relevant behavioural capacities are realised, but also with the claim that nonrepresentational explanations are more appropriate. Before I get to this, I will examine a more straightforwardly representationalist option, involving the claim that Sensorimotor Understanding is Knowing What Would Happen If You Did Act, and depends on the deployment of a particular kind of internal representation.

4.3 - Sensorimotor Understanding as Subpersonal Representation

Suppose, in accordance with some of Noë's (2004) claims, that possessing and exercising Sensorimotor Understanding neither requires bodily movement nor an ability or disposition to move. To adopt this position is, by definition, to reject a number of claims that could be deemed behaviourist. In addition, suppose that action is not what Sensorimotor Understanding is primarily for. Instead, let's take it that Sensorimotor Understanding is a kind of knowledge about the P-properties that presently obtain, and the P-properties that would obtain given various possible movements. As I put it earlier, Sensorimotor Understanding, on this view, has primarily a belief-like direction of fit. Sensorimotor Understanding continues to have one non-trivial conceptual link with bodily behaviour, namely that it incorporates knowledge about the consequences of possible bodily movements.

Noë (2010) continues to deny that conscious vision is primarily for action, but argues that if action-guidance is one purpose of conscious vision, then conscious vision requires Sensorimotor Understanding. He contrasts the sensorimotor approach, as usual, with the claim that perception consists of the activation of an inner world model. On the inner model conception, perceiving involves knowing what objects are out there, but not what bodily relation you stand in to them. Lest one think that Noë is targeting only a pictorial theory which no one endorses, notice that any representationalist theory is in this position by default if the representational content it endorses does not somehow specify a bodily relation. Noë claims that if your perceptual awareness did not represent your bodily relation to the world, then the world would not seem phenomenally 'present' to you. Even if this is not true, it is not apparent how your perceptual awareness in this event could help guide action. Noë observes that it must be possible to integrate what you see with what you do, and argues that it is a

virtue of the sensorimotor theory that it shows how this integration is possible.

An upshot of this point, we might think, is that Sensorimotor Understanding necessarily involves an ability to make skillful bodily movements, so long as something else does not block the ability - for example paralysis, or optic ataxia, an inability to co-ordinate muscle movements. This is one way of weakly motivating the conceptual link between perception and action implied by the title of Noë's (2004) book. It also justifies the description of Sensorimotor Understanding as practical, but only if we take 'practical' to mean knowing-that *and* knowing-how, as Stanley and Williamson do and Noë denies. Sensorimotor Understanding, on the present view, is knowing what would happen if you moved, a knowing-that, accompanied by a knowing-how to act, which depends on the knowing-that. However, the present link between perception and action is less tight than the one expressed by Noë's claim that "*What we perceive* is determined by *what we do* (or what we know how to do)" (2004, p. 1), since the present view is in essence only signed up to the more pedestrian claim that what we do or know how to do is determined by what we perceive.

Contradicting Noë's (2004) discussion of practical knowledge, Sensorimotor Understanding on the present view is not merely a practical ability, but a belief-like understanding of SMCs that enables behaviour. This means that the sensorimotor approach owes an explanation of what Sensorimotor Understanding is and, in particular, how it is physically constituted or realised. Since it is belief-like, it appears to require truth-evaluable content. Despite asserting that we cannot usually express our understanding of SMCs using propositions or bring it to bear in "explicit deliberative judgement" (p. 187), Noë asserts that Sensorimotor Understanding is conceptual, on what he takes to be a legitimately unrestricted account of what conceptual thought requires. He changes this position in later work, asserting instead that Sensorimotor Understanding involves a kind of nonconceptual content, but maintaining that Sensorimotor Understanding occurs on a continuum with conceptual skills rather than belonging to a radically different kind (Noë, 2012). This does better justice to the suggestion that Sensorimotor Understanding involves content which the perceiver cannot fully express using propositions, even, presumably, by pointing to an object and saying '*that* sensorimotor law'. This raises the possibility that Sensorimotor Understanding supervenes on the deployment of neural representations bearing nonconceptual content.

Seth (2014) has recently offered a specific representationalist proposal intended to do justice to some aspects of the sensorimotor theory. His account rejects the sensorimotor theory's explicitly enactivist apparatus, namely anti-representationalism, vehicle externalism and the endorsement of dynamical systems explanations. However, he endorses the account's approach to naturalising phenomenal qualities, and, in particular, Noë's (2004) insights into visual phenomenology, including presence-in-absence. Seth combines these aspects of the sensorimotor theory with predictive coding (see, Clark, 2013; Hohwy, 2013). Like the orthodox approach, predictive coding accounts of perception hold that the brain constructs internal representations designed to represent features present in the outside environment as accurately as possible, based on the limited information available at one time via the sense organs. Orthodox theories make sensory inputs the starting point, and propose that the brain uses them to build up an inner model. By contrast, predictive coding accounts hold that the brain constructs models of the environment which it 'predicts', using Bayesian inference, will prove to be in keeping with the worldly state of affairs that really obtains.

The difference between this and the orthodox approach is that sensory inputs here serve as error correction signals; instead of acting as precursors to the construction of an inner model, the sense inputs function to tell the brain when the inner models it has already deployed are wrong. Further, the predictive approach holds that the brain simultaneously holds multiple inner models organised hierarchically into tiers at increasing levels of abstraction. Only the lowest model in the hierarchy uses sense inputs themselves as error correction signals. This model functions as an error correcting signal for the model above, and so forth. An example of a substantive explanatory difference between orthodox and predictive coding accounts occurs in the phenomenon of *repetition suppression*, a tendency of the brain to respond less to a stimulus as it becomes more familiar. Seth cites work showing that repetition suppression fails to occur when the subject does not expect the stimulus to be repeated (Summerfield, Trittschuh, Monti, Mesulam and Egner, 2008). He observes that this is intelligible within a predictive coding approach, since the approach makes expectations a central feature, whereas there is no natural way to make this intelligible within an orthodox framework, where sense inputs come before the models which make use of them.

Seth endorses some of the Noë's (2004; 2012) insights about visual phenomenology.

One example of this, which I have discussed already, concerns shape-related P-properties and objective properties. You are confronted with a coin, tilted at an angle, which is therefore elliptical in aspect. Thanks to your Sensorimotor Understanding, you appreciate both that it has an elliptical aspect and that objectively it is circular. It is a key virtue of the sensorimotor approach that it explains how you are sensitive to objective properties while at the same time, necessarily, being sensitive to them from some perspective rather than from nowhere. Another aspect of this phenomenon is ‘presence-in-absence’, the tendency of environmental features which are strictly speaking out of view to show up in visual consciousness anyway. Noë’s (2004) prototypic example concerns looking at a tomato. In one sense, you can never see the whole tomato at once, but only one side of it, while the other side remains out of view. According to Noë’s phenomenology, which Seth endorses, the back of the tomato nonetheless has presence-in-absence, which is to say your visual experience is of a whole tomato rather than something like a 2D cut-out.

Seth argues that a predictive coding account can do justice to this, proposing that the predictive models employed by the brain in perception incorporate not only predictions about the features of the present environment causally responsible for present sense inputs, but also counterfactual predictions about the changes in sense input that would result from a range of possible movements. Perceptual presence is explained by these counterfactual predictions, and the more predictions there are about movement-related counterfactuals relating to some feature, the greater the degree of felt phenomenal presence there is associated with that feature.

For my present purpose, the crucial feature of Seth’s account is that it provides a better developed explanation of the mechanism that underlies Sensorimotor Understanding than found in the original statements of the sensorimotor approach, and does justice to the ‘belief-like’ construal of Sensorimotor Understanding I have been addressing in this section. It is notable that adopting this view means adopting a very conservative rendering of the sensorimotor account. It means endorsing representationalism, rejecting dynamical entanglement and ECM, and by focusing on internal models, downplaying the role played by embodied interaction even as causal scaffolding.

Sensorimotor mastery, and in particular Sensorimotor Understanding, one of the

most important notions in the sensorimotor theory, is ambiguously defined by the theory's original statements. At times, it appears to be implicit knowledge of the SMCs that presently obtain, and nothing more. In this case, the obvious way to explain Sensorimotor Understanding is in keeping with orthodox cognitive science via a representationalist story. Seth's account has the virtue of doing justice to Noë's (2004) insights about visual phenomenology and perceptual presence, and of allying this aspect of the sensorimotor theory with insights from the currently popular predictive coding approach. If nothing else, Seth's account also has the virtue of beginning to repair the sensorimotor theory's serious explanatory deficit by accounting for the precise mechanism that realises Sensorimotor Understanding. However, this comes at a cost. For anyone who is motivated, on whatever independent grounds, to endorse the enactivist theses of anti-representationalism, vehicle externalism and dynamical systems, Seth's account will not help, since it serves only to co-opt insights from the sensorimotor theory into a much more conservative representationalist framework. In particular, accounting naturalistically for representational content brings with it a heavy and as yet unresolved explanatory burden (Hutto and Myin, 2012), as I discuss further in chapter 6.

Moreover, the approach has costs for the sensorimotor theory itself. If we endorse internalism and representationalism, then the sensorimotor theory no longer gets to make distinctive use of the appealing proposition that the brain offloads the burden of modelling onto the environment itself. Although any variety of representationalism can observe that the amount represented at any one time is minimal, and that as active perceivers we move around to acquire new information, the implication of Noë's (2004) account is that Sensorimotor Understanding is computationally more efficient than inner modelling, because we merely need to be poised to interact with the environment, rather than modelling the environment and *then* becoming poised to interact with it on the basis of this model.¹⁵ Seth's proposal also has the drawback, from the perspective of sensorimotor theory, of making the sensorimotor account hostage to phenomenological fortune, hanging as it does in large part on Noë's

¹⁵ Of course, whether the non-representationalist version of the sensorimotor account really implies a more efficient process than a representationalist alternative depends on how the representationalist and non-representationalist alternatives are precisely cashed out.

description of visual phenomenology. Clark (2008), for instance, suggests that the feeling of presence-in-absence could be well accounted for by the general feeling that there is more to be seen, rather than any more explicit representation of or knowledge about the sensory consequences of possible movements. For these reasons, I endorse a different view of Sensorimotor Understanding, which I lay out in the next section.

4.4 - Sensorimotor Understanding in Bodily Capacities

Notwithstanding suggestions by various commentators that the sensorimotor account requires representation, the theory is, in at least one sense, strongly anti-representationalist. The idea is that there is no need to represent the world, because it is already present to you. This is reflected by the outside memory thesis (O'Regan, 1992; O&N), and in Noë (2004; 2012) by the idea of presence-as-access, which implies that perception does not depend on representing even a small part of the outside world. One way of thinking about the theory's anti-representationalism is to think of it as a constitutive claim about perception at the personal level. On this proposal, we are not, as persons, in touch with representations of the world, but stand in an unmediated relation to the world itself. This is not necessarily incompatible with the claim that perception depends on internal representations subpersonally, a point I return to later.

The sensorimotor account's way of cashing out this (personal level) relation with the outside world is to say that it involves possessing and exercising a set of bodily capacities; specifically, capacities to skillfully interact with the outside environment in ways that draw on an implicit understanding of the SMCs that obtain. Consider that on a representationalist account, perception involves representing the world, and interacting with the environment in a goal-directed way is a contingent consequence of the fact you perceive it. Since the role of the 'representation' is played in the sensorimotor account by the world itself, it follows that perceiving itself in the sensorimotor account is equivalent to possessing and exercising the capacity to skillfully interact with the world in a goal-directed way. This provides a solid motivation for the claim that perception depends on action and Being Able to Act, which, as I said earlier, implies a desire-like direction of fit.

The question remains, however, of how we account for Sensorimotor Understanding in its guise as a belief-like faculty, a Knowing What Would Happen If You Did Act. The best approach, I suggest, is to avoid reifying Sensorimotor Understanding. We should not think of it as the sort of thing that depends, by definition, on a particular sort of concrete entity, for example a conscious or purely subpersonal representation, which causally mediates behaviour. Instead, we should think of it as a set of necessary conditions that our capacities for practically-oriented, goal-directed behaviour must meet to implicate perception and perceptual consciousness.

To understand the explanatory relation between Sensorimotor Understanding and behaviour, it will help to distinguish between some different varieties of explanatory relation theoretical entities can stand in to one another. Bermudez (2000) distinguishes between ‘vertical’ and ‘horizontal’ explanations in cognitive science (not to be confused with Hurley’s 1998 usage of the vertical/horizontal metaphor). Horizontal explanations exist at one level, the personal or subpersonal, and describe temporally distributed events. So if we are offering a horizontal explanation and we say that A causes B, we mean that A happens and this causes B to happen a short time later, where A and B occur at the same explanatory level. On Seth’s (2014) proposal, we might imagine that neural representations of the sensory results of possible movements send information to motor output areas of the brain, and so generate behaviour. This would be a horizontal explanation. Vertical explanations describe events that happen simultaneously at different explanatory levels. In this case, A might be a subpersonal computation, and B might be a psychological state possessed by the whole person. People sometimes use the word ‘cause’ in vertical explanations too, although it is not causation of the sort that implies temporal extension. We might also say, here, that A *enables* B. If we said that the deployment of the subpersonal representations described by Seth enable perceptual consciousness, this would be a vertical explanation.

The important thing, for the moment, is that the explanatory relation between Sensorimotor Understanding and a capacity for skillful behaviour is neither a vertical nor a horizontal one. Sensorimotor Understanding, I propose, is a logical condition that the behaviour enabled by our behavioural capacities must meet in order to constitute perception and perceptual consciousness. The identity between perception and skillful behaviour arises

from the fact that to perceive is not to model the world and then use this model, but merely use the world in the way traditional theories hold that we use an inner model. The fact that the behaviour needs to manifest a belief-like Sensorimotor Understanding is an upshot of the straightforward claim that it would be impossible to skillfully interact with the world in more than a very basic way if you did not have the ability to anticipate, in at least some cases, the sensory results of your movements. On this basis, behavioural capacities manifesting Sensorimotor Understanding can be described as a condition of possibility for perception. This is one way of understanding the claim that Sensorimotor Understanding *grounds* perception-implicating behaviour.

The view I am advocating is neutral with regard to whether or not Sensorimotor Understanding is enabled by the deployment of subpersonal representations. However, it does not require internal representation as a matter of definition or logical necessity. Sensorimotor Understanding should be thought of as a characteristic of behaviour; an adverb or adjective used to describe behaviour or a behavioural capacity, and not a noun. Although there is no adverb ‘understandingly’, we might say that a behaviour ‘shows understanding’, and mean that the understanding is constituted by the behaviour rather than that the behaviour merely demonstrates the existence of a pre-existing entity called understanding. Since Sensorimotor Understanding is a logical condition for certain kinds of skillful behaviour, appropriate behaviour is almost sufficient by itself to guarantee the truth of an ascription of Sensorimotor Understanding. We should add to this an extra requirement that the brain makes an intelligible contribution to the behaviour, since it is doubtful that a person in a metaphysically possible world who exhibits appropriate behaviour but lacks internal states to enable the behaviour could be conscious.

It may help to compare the current view of Sensorimotor Understanding with Dennett’s (1987) intentional stance, which I discussed in chapter 2. According to the intentional stance, mental states such as beliefs and desires exist just as long as ascribing them to a person or other agent helps an outside observer explain the person’s behaviour. Like Sensorimotor Understanding, this implies the existence of a set of necessary and sufficient conditions that an agent’s behavioural repertoire must meet for the agent to possess a particular mental state. Sensorimotor Understanding is correctly ascribed to any perceiver

whose behavioural repertoire meets the appropriate conditions. In contrast with a certain understanding of the intentional stance, however, Sensorimotor Understanding is, in a robust sense, real, and not observer-relative. This is because certain kinds of skillful behaviour require, as a matter of necessity, a sensitivity to certain presently obtaining SMCs.

If we can account physicalistically for the relevant behavioural capacities without appeal to neural representations, this will be enough to account for perception, and the belief-like Sensorimotor Understanding will come for free. The suggestion that it is possible to explain the relevant behavioural capacities without appeal to representations is controversial, of course, and nothing I have said so far shows this can be done. However, it is important that Sensorimotor Understanding, on the view I propose, does not require subpersonal representation by conceptual necessity. Moreover, the present view is compatible with the possibility that perception does involve neural representations, but not neural representations of SMCs. This holds open the possibility of a subpersonal representationalist explanation of Sensorimotor Understanding quite different from Seth's approach.

We need, of course, to know something about the behaviour required to justify ascriptions of Sensorimotor Understanding. In the case of P-properties like relative location, P-size and P-shape, Sensorimotor Understanding is manifest in the capacity to make appropriate movements in sensorimotor space. For instance, an understanding that something is to the left might be manifest in the capacities to walk, run, crawl (etc.) in a leftward direction for the purpose of bringing the object into view. When we experience something as being to the left, this usually implicates, as Evans (1985) says, a capacity to walk *or* run *or* crawl, or perform any action from an indefinitely wide range of possible actions in order to reach the object. If we know that someone intends to face the object head-on, and we see them walking toward it, we could, in a manner reminiscent of the intentional stance, interpret them as possessing an implicit understanding of the fact that this movement will bring the object into view, hence as possessing Sensorimotor Understanding. However, ascribing a perceptual state to a person or animal on the basis of this single behaviour would be ascribing too much. The behaviour could be a reflexive response to a stimulus, where perception ought to enable an ability to respond flexibly using a wide repertoire of behaviours. So the subject ought to have the ability to run, crawl, etc., as well as walk in order to reach the stimulus; the capacity

has, in this sense, to admit some degree of generalisability. We can do justice to Noë's contention that the experience of objective properties is grounded by the same faculty of Sensorimotor Understanding as the experience of P-property if we cast the experience of objective properties as involving the exercise of a behavioural capacity, too. In common with Sensorimotor Attunement, appealing to this capacity differs from the behaviourist appeal to input-output relations, because it does not specify which outputs should follow when particular inputs obtain, but which outputs should follow when particular SMCs obtain.

There are various similarities and differences between Sensorimotor Understanding and the notion of Sensorimotor Attunement, which I reconstructed from O&N in section 3.3.1. In Noë (2004), Sensorimotor Understanding is largely if not entirely framed as a personal-level capacity, meaning that its operations are accessible to the conscious subject. This reflects the fact that we are conscious of possessing and exercising the relevant understanding, or at least we always have the possibility of becoming conscious of it, even if there are cases in which we are not conscious of our behaviour because we are sufficiently absorbed. Evans (1985) claims that to experience an object as occupying a particular spatial location, the conscious subject must understand the movements required to get there. This might be equivalent to nothing more than being aware of what your bodily dispositions are. If you are aware of the movements you must make to reach an object that is to your left, for instance, then you are aware of the SMCs associated with the object's being on the left, i.e. the movements that will result in your facing the object head-on. In this sense, the SMCs involved in 'being to the left' are transparent to you (i.e. you know what they are). The experience of P-colour, the colour a surface presents when viewed from a particular location and under a particular lighting condition, is a different sort of case. From the perspective of a conscious perceiver, there is no transparent connection between this aspect of colour phenomenology and possibilities for bodily movement. Noë (2004) thus suggests that P-colour experience is accounted for by a subpersonal (for which we can read subdoxastic) variety of sensorimotor mastery (p. 140). Since I am attempting to offer a maximally coherent reconstruction of the sensorimotor theory, I will take it that this is accounted for by Sensorimotor Attunement.

This brings to the fore an important difference between Sensorimotor Attunement and Sensorimotor Understanding. When I attempted to do justice to O&N's notion of

Sensorimotor Attunement, I noted that there was more than one way the theory could propose we identify the specific SMCs to which a perceiver is presently attuned when she exercises the relevant behavioural capacity. I suggested that one option is to endorse a Dennett-style interpretative stance and identify the SMCs by appeal to a combination of teleology or proximal goal and present behaviour. I went on to endorse a more austere alternative in which we suppose that a goal-directed agent masters a given set of SMCs if she responds differentially to them in her behaviour, regardless of whether or not one could identify what the SMCs are by running the interpretative stance.

In this event, the agent does not have to know or represent what the SMCs are, she merely has to discriminate them. This is compatible with the suggestion that conscious perceivers do not have cognitive access to the SMCs involved in P-colour. This contrasts, however, with the conscious Sensorimotor Understanding that applies to P-shape, location and movement. In this case, where we do have the option of running a behavioural test to see if a subject possesses Sensorimotor Understanding, the subject in some sense consciously knows what the SMCs are, because the SMCs themselves, and not just discrimination of the SMCs, are identifiable from the subject's behavioural repertoire - e.g., she is going to the left, she wants to face the object head on, so she has mastered the fact that going to the left will result in facing the object head on.

To conclude this section, I will attend to some interesting remarks by Glock (2008) on the subject of animal thought. In the piece, Glock attempts to tread a middle path between linguicism, which claims that non-human animals lack thought altogether because they lack language, and the mentalist claim that animal thought differs from human thought only in degree and not kind. Although this is a quite different topic to the one that concerns us here, it chimes with Noë's (2012) stated ambition of accounting for continuities between thought in humans and other animals, and between non-linguistic thought and sensorimotor skills in mature humans. Glock criticises Searle (1994), who claims that the best way to identify whether or not an animal is conscious is to identify neurophysiological features correlated with consciousness and check whether the animal possesses them or not. Glock points out that this only works if you have a way to test outwardly observable behaviour for consciousness in the first place. Even once you have identified the relevant features in one

species, there is no guarantee that another species which also possesses those features is also conscious purely because it also possesses them. Glock's solution to the question of animal minds is to focus on capacities that can be manifested in behaviour instead. The capacities in question are not mere dispositions to discriminate one object from another, but more flexible abilities, for example, to respond to distinct objects differently in one situation but not another, and to correct mistakes. The idea is not that thoughts reduce to certain behavioural dispositions, but that they ought only to be ascribed on the basis of them.

In the case of human perceptual consciousness, the problem is different. There is no need to concern ourselves with the question of whether or not other humans enjoy perceptual consciousness. But a question does arise with regard to what perception consists of. It is possible that the only thing which would licence us to claim that a particular neurophysiological state gives rise to perceptual consciousness is its being correlated with a behaviour too. As I have suggested here, that is behaviour meeting a certain set of conditions described as Sensorimotor Understanding. This is an epistemic motivation for ascribing perceptual skills on the basis of behavioural capacities.

However, there is a deeper motivation, which is the possibility that perception just *is* a capacity for skillful bodily interaction with the environment. There is no default reason for thinking this claim fares any worse than the claim that perception is constituted by the activation of an inner world model. If the capacity claim is true, we get the benefit of a direct realist account of phenomenal qualities, which I mentioned in section 3.6 in the discussion of ECM, as well as the epistemic benefit of a non-mysterious way to explain how physical states and processes give rise to perceptual consciousness. In the next chapter, I claim that sensorimotor theory can enjoy partial success as an account of phenomenal qualities even if the present view fails. In chapter 6, I defend the claim that perception is constituted by the possession and exercise of a set of behavioural capacities.

5 - Phenomenal Character: A Sensorimotor-Representation Hybrid

The next two chapters address a response made by Clark (2008) to the sensorimotor account in which he claims that perception plausibly consists of the construction and deployment of representations standing in for types, categories and relative locations, and not the exercise of sensorimotor skill. In chapter 6, I pursue some more radical claims arising from the view of Sensorimotor Understanding as a set of conditions that behaviour must meet, and the attendant possibility that perception does not involve representation at all. In this chapter, however, I explore a response by the sensorimotor theory to a worst case scenario in which it is established that perception does depend on representation, and not on behavioural capacities. I argue that even in this event, the sensorimotor theory has the ability to provide a tenable, compelling account of phenomenal qualities. This chapter is structured as follows. In the first section, I outline a criticism levelled by Clark (2008) against the sensorimotor theory. In the second section, I propose a line of response that could be adopted by a conservative rendering of the sensorimotor approach. In the third section, I examine an apparent problem for this response, and explain how it can be overcome. In the fourth section, I offer further reasons to endorse the sensorimotor account of phenomenal qualities. In the fifth section, I examine the relation between sensorimotor theory and Rupert's (2011) model of the 'massively representational mind'.

5.1 - The Non-Sensorimotor, Skill-Based Approach

In a critical response to the sensorimotor theory, Clark (2008) endorses, in broad outline, an approach to perceptual consciousness which he suggests is both more plausible than the sensorimotor account, and incompatible with its key claims. Clark's preferred approach in part comprises a philosophical approach to phenomenal qualities or 'qualia', which he argues are properly accounted for by the perceiver's exercise of 'epistemic' skills and not the sensorimotor skills that Noë and O'Regan appeal to. The approach also comprises an empirically-oriented account of the cognitive processing that underlies perceptual consciousness. Drawing in particular on Milner and Goodale's work on dorsal and ventral

stream perception (e.g., Milner and Goodale, 1995), Clark suggests that the empirical evidence demonstrates that conscious perception is not primarily the result of a sensorimotor process. I will begin by outlining Clark's preferred approach to phenomenal qualities, and subsequently move on to his claims about cognitive processing.

Skill-based accounts of phenomenal quality propose a way to account naturalistically for the phenomenal qualities featured in perceptual consciousness, for example the look of red. On this approach, a given phenomenal quality is accounted for by nothing more or less than the subject's ability to skillfully discriminate, in some relevant way, a feature of the environment associated with the quality from other features they have the ability to consciously perceive. One such account is offered by Pettit (2003). As Pettit notes, traditional accounts imply that a subject's capacity to discriminate one property from another is merely the contingent outcome of a difference in phenomenal look. Conversely, skill-based accounts claim that there is nothing more to experiencing a given phenomenal look than exercising the associated discriminatory capacity. In Pettit's examples, you exercise these capacities in making judgements, for instance that an object is red, or a brighter shade of red than some other object, as well as in bodily behaviour, for example reaching out appropriately to catch a fast-moving ball. An object looking a particular shade of red to you is, on this view, no more than it looking to you like you can discriminate it in an appropriate way from other shades of red, and other colours, and so forth. Presuming that there is no fundamental obstacle to describing the exercise of these capacities for discrimination in physicalistic vocabulary, there is no deep-seated obstacle on this approach to understanding phenomenal qualities physicalistically.

Clark takes it that the sensorimotor theory is itself a skill-based approach. On most ways of understanding the account this is true, although we must take care to be clear on why this is so. The sensorimotor approach proposes that the look of red, like other phenomenal qualities, is properly accounted for by a sensorimotor law, which is a systematic way that sense input signals originating from a relevant object change in line with movements and varying ambient lighting conditions (e.g., Noë, 2004). One way to understand the concept of sensorimotor law, therefore, is to think of it as something that can be described solely by reference to counterfactual conditionals of the form 'if the subject moved like this, sense input

would change like this' (suggested by Noë, 2004). So understood, sensorimotor laws comprise what Burhmann, Di Paolo and Barandiaran (2013) describe as the subject's *sensorimotor environment*. These laws have a motor output to sense input direction of fit, since they characterise how sense input will change in line with given movements, but do not describe how the subject will move given particular sense inputs. So understood, sensorimotor laws may, in principle, be sufficient to explain degrees of similarity and difference between specific phenomenal qualities where it is assumed that the subject is in some appropriate way sensitive to those sensorimotor laws. But they could not explain how a perceiver can be in conscious states characterised by those qualities, considering that the relevant counterfactual conditionals might be true even if the perceiver was dead, or their brain were made of jelly (etc.).

To explain how perceivers do come to be in the relevant conscious states, the sensorimotor approach posits that in some way or another we master sensorimotor laws. Depending on how we construe the notion, mastery of the sensorimotor environment could include internal representations of counterfactual conditionals, perhaps even propositional knowledge, without necessarily calling on any bodily skills whatsoever. As it happens, however, most versions of the sensorimotor account individuate phenomenal qualities by reference only to patterns of counterfactual dependence that the perceiver manifests a sensitivity to in their bodily movements.¹⁶ Burhmann, Di Paolo and Barandiaran (2013) characterise the perceiver's repertoire of bodily responses in a given sensorimotor environment through what they term a *sensorimotor habitat*, which this time comprises the movements the perceiver makes in line with given sense inputs and the changes in sense input that follow. The need to refer to a perceiver's sensorimotor habitat and not just their sensorimotor environment is evident in O&N and O'Regan (2011), regardless of their further appeal to higher-order

¹⁶ Noë (2010) actually denies that the sensorimotor theory is a skill-based approach, on the ground that perception depends on understanding SMCs rather than merely being able to skillfully discriminate them. However, the frequent invocation of sensorimotor 'skill' invites a reading in which the sensorimotor theory is a skill-based approach. This reflects the tension between competing directions of fit I discussed in section 4.2. Since I concluded that Sensorimotor Understanding ought to be thought of as depending on knowing how to act as well as knowing what would happen given possible movements, I think it appropriate to consider the sensorimotor theory a skill-based approach.

tracking of those bodily interactions. Noë (2004) does not say anything explicit to suggest that a perceiver can only be conscious of sensorimotor laws featured in their sensorimotor habitat. However, the claim is not explicitly retracted in Noë's (2004) solo account. At times, Noë attempts to avoid endorsing an unduly strong link between perception and action, and in particular denies that you necessarily have to be moving your body right now in order to see. This does not compromise the suggestion that you can only have perceptual experiences characterised by SMCs featured in your sensorimotor habitat, so long as we take the sensorimotor habitat only to characterise ways you are *disposed* to skillfully interact.

Clark agrees that phenomenal qualities should be explained by reference to the exercise of skills. However, instead of appealing to bodily responsiveness to patterns of sensorimotor dependence, his favoured account appeals to personal-level responsiveness in reasoning and action planning to categories, types and locations. He takes it that 'epistemic' skills of this sort either do not call upon capacities for bodily movement, or at least do not do solely call upon such capacities, and are distinct from the bodily skills appealed to by the sensorimotor theory. To keep things unambiguous, I will describe epistemic skills, as Clark understands them, as 'non-sensorimotor' skills, leaving the designator 'epistemic' neutral on whether or not the skills in question are identical to capacities for bodily movement. Importantly, Clark argues that sensorimotor and non-sensorimotor skill-based accounts of phenomenal qualities share a common explanatory virtue when it comes to accounting for phenomenal qualities non-mysteriously; indeed, nothing he says is intended to call into doubt the view that the sensorimotor and non-sensorimotor skill-based accounts are, in principle, equally well-suited for resisting the claim that phenomenal qualities cannot be accounted for naturalistically. This makes it plausible that Clark thinks the appeal to skills is enough to solve this problem without even getting into the question of whether they are sensorimotor or not. His objection to sensorimotor theory, by contrast, is that it provides an implausible account of the capacities that are, as a matter of fact, deployed in perceptual consciousness.

Clark's objection partly reflects a complaint that the sensorimotor theory is untenably committed to supposing that the range of phenomenal qualities featured in an agent's perceptual awareness is no less fine-grained than the discriminations evident in the agent's most finely sensitive bodily behaviour. This commitment is a consequence of the

sensorimotor theory's claim that two perceivers could only share identical perceptual phenomenology if they are identically embodied. O&N, for example, say: "For two systems to have the same knowledge of sensorimotor contingencies *all the way down*, they will have to have bodies that are identical *all the way down*" (p. 1015), with a similar claim appearing in Noë (2004).

In fact, it is not apparent that the sensorimotor theory has any essential need to make the claim. Logically-speaking, the sensorimotor version of the skill-based view only needs to say that for every difference in phenomenal quality there is a corresponding difference in the bidirectional patterns of dependence between input and output featured in the perceiver's skill-driven bodily interactions. It need not make the inverse claim that every discrimination manifest in an agent's behaviour must be reflected in the phenomenal character of their perceptual consciousness. The sensorimotor theory faces pressure to say that the input-output patterns relevant to perceptual experience are very fine-grained in order to avoid an objection levelled by Block (2001). Block's charge is that the sensorimotor theory must claim that a rudimentary machine, differently embodied but manifesting comparable input-output relations, undergoes human-style perceptual experience. O&N's reply is that sameness of perceptual experience depends on sameness in fine-grained input-output relations, which themselves depend on being embodied in a particular way. However, they need not have claimed that sameness of embodiment is required *all the way down*. The sensorimotor theory could respond to the charge just as well by claiming that two perceivers must be embodied in a sufficiently *similar* way in order to enjoy identical perceptual experiences.

If Clark's view were merely that we should account for qualia by appealing to coarse-grained bodily skills, it would not be inconsistent with the most indispensable commitments of the sensorimotor theory. More seriously, however, his comments imply that the relevant capacities may not be bodily capacities at all. The capacity to reason, certainly, does not on the face of it appear to be a bodily skill, and nor does the capacity to plan an action if this is taken to mean choosing an action rather than actually performing it. Clark bases his response largely on empirical work, which he alleges shows that perceptual consciousness is not realised by sensorimotor coupling but a process involving the deployment of internal representations that are in large part unaffected by occurrences at the sensorimotor

periphery.

In particular, he cites Milner and Goodale's (1995) dual systems account of vision. According to the dual systems model, the brain's dorsal stream is associated with a visual system responsible for fine-grained sensorimotor coupling but not conscious visual awareness. Clark does not raise any objection to describing this bodily behaviour as the exercise of sensorimotor mastery, and we can plausibly take this activity to be identical to the activity characterised as the lowest tier of sensorimotor mastery in O'Regan's (2011) multi-tier sensorimotor account. The other visual system is realised primarily by activity in the ventral stream, and is alleged by Clark to enable coarser-grained "reasoning, choice and action selection" (Clark, 2008, p. 190), a description which is strikingly akin to the "planning, decisions, or other rational behaviour" (O'Regan, 2011, p. 121) featured in the second of O'Regan's four tiers. Clark argues that ventral stream activity is not sensorimotor mastery but involves the deployment of representations that, in line with the epistemic skill-based approach to qualia, stand-in for attributes such as "categories, types, and relative locations" (p. 192), optimised for the exercise of epistemic skills, i.e. to enable to person to "sift, sort, select, identify, compare, recall, imagine and reason" (p. 192).

In this vein, Clark says: "[T]he best functional and architectural explanation, according to Milner, Goodale, and others, is that conscious perceptual experience reflects the activation of representations that have less to do with the fine details of world-engaging sensorimotor loops and more to do with the need to assign inputs to categories, types, and relative locations" (p. 192). The role of representation here needs to be treated with care. O'Regan (2011) explicitly denies that conscious states can be adequately accounted for by reference to the 'activation' of representations, a locution featured in the Clark quote above, on the ground that accounts of this sort cannot adequately explain why the conscious states in question have one phenomenal character rather than another. Clark appears to endorse O'Regan's view, at least insofar as he endorses a skill-based account of phenomenal qualities and not, for instance, an account that says phenomenal qualities are accounted for directly by the content of the purported representations (cf., Tye, 2000). Being mindful of this, Clark's view ought to agree that it is not *just* the 'activation' of a representation, but the contribution the representation thereby makes to the exercise of the relevant personal-level skillful activities

like sorting and sifting that contributes to an explanation of phenomenal qualities. At the subpersonal explanatory level, it is an open question how a given ventral stream representation contributes to the exercise of personal-level skill. An explanation here may appeal to some combination of its content, its functional role within a given system, and that system's functional role in the wider cognitive economy. Whichever features an endorser of this approach appeals to, they must explain what contribution these features make to the exercise of a personal-level ability to skillfully discriminate aspects of the visual environment.

In the discussion that follows, I will describe the account of perception and phenomenal qualities that Clark endorses as the *non-sensorimotor, skill-based approach*. The approach is framed by Clark as a chapter-long response to the sensorimotor theory rather than a developed account of perception, but in broad terms it advocates a number of positive claims about visual processing and the phenomenal qualities that characterise visual consciousness. The account's claims can be summarised as follows: (i) There are two visual systems, which operate in the way described above; (ii) Ventral stream vision deploys representations that represent properties like type and location rather than sensorimotor law; (iii) The ventral stream deploys these representations in the course of enabling the exercise of skillful capacities like sorting, sifting, classifying, selecting and others, which for shorthand I will refer to just as sorting and sifting; (iv) The properties sorted and sifted are not sensorimotor laws, but types, categories and locations; (v) The appeal to sorting and sifting offers, in principle, a non-mysterious, physicalistic account of the phenomenal qualities associated with conscious perception; (vi) Sorting and sifting are not sensorimotor skills. They are realised by subpersonal systems which are, to a significant degree, informationally encapsulated from the operation of the dorsal stream and motor output systems. This enables them also to be largely insensitive to the fine-grained morphological characteristics of the agent's body.

5.2 - The Policy of Containment: A Conservative Response

Clark (2008) suggests that the sensorimotor approach to phenomenal qualities enters the running alongside competing skill-based accounts that share the explanatory benefit gained by appealing to skills, but lack what he regards as the undue insistence that the skills in question

are necessarily sensorimotor. His response to the sensorimotor account in that chapter is not primarily concerned with phenomenal qualities, however, but with the character of the cognitive processing that realises perceptual consciousness. He argues that the empirical evidence concerning ventral stream perception challenges the sensorimotor account by indicating that perceptual consciousness supervenes on internal states, in particular representations, that stand-in for categories, types and locations rather than sensorimotor contingencies.

To respond, the sensorimotor theory could adopt an approach I will call the Policy of Containment. The response aims to contain Clark's criticisms by restricting the explanatory role played by law-governed sensorimotor interactions to one domain, while asserting that the evidential force of the empirical work Clark cites is restricted to another. Specifically, it contends that the appeal to sensorimotor coupling is not meant to characterise cognition in humans, but solely to account for the phenomenal qualities that feature in perceptual consciousness. As an account of qualia, the sensorimotor approach enters one classic debate, prompted by philosophers like Jackson (1982) and Chalmers (1996), over whether phenomenal qualities are an intelligible part of the physical world. Assuming that phenomenal qualities are physical, it also enters mainstream debates over what sort of physical qualities they are, included in which is the question of which characteristics make one phenomenal quality different from another. On this latter set of questions, the approach enters the field alongside the accounts offered, for example, by Tye (2000), who proposes that phenomenal qualities reduce to representational content; Block (1996), who suggests they reduce to some non-representational neural property; and Pettit (2003), who, in common with sensorimotor theory, offers a skill-based account.

It is important to emphasise that it is not obvious, pretheoretically at least, that an account can address the above topics simultaneous to characterising the cognitive processing that underlies perceptual consciousness. Tye addresses cognition and phenomenal qualities at once, since he proposes that phenomenal qualities are individuated and reduced to the physical by reference to the content of neurally-realised representations, and thereby overtly commits himself to the existence of internal representations in cognition. Clark's account also appears to suggest that there is a one-to-one mapping between the properties represented

subpersonally by the ventral stream and the properties skillfully discriminated in conscious perception. Despite this apparent mapping, however, Clark appeals to quite different things respectively to explain phenomenal qualities and cognition. In the former case, it is the exercise of personal-level skills; in the latter, the activation of representations. In this respect, Clark's account is quite different to Tye's.

On the gloss endorsed by the Policy of Containment, the sensorimotor theory is open to endorsing an account of cognition along the lines of the one advocated by Clark (2008), in particular the claim that visual consciousness is realised by representations that stand in for categories and types. In this respect, the theory endorses a traditional representational approach, although it emphasises that cognitive processing is importantly scaffolded by bodily exploration, a claim which does not conflict with Clark's characterisation of ventral stream vision. To account for phenomenal qualities, however, the sensorimotor theory suggests that we need only focus on regularities in the agent's skill-driven sensorimotor coupling, avoiding altogether the more opaque territory of the brain. On this reading, the theory can be regarded as a kind of *sui generis* hybrid between traditional and radically embodied accounts of mind. The success of this approach depends on showing that there is no incompatibility between its respective accounts of phenomenal qualities and cognitive processing, and, additionally, that there are positive reasons to endorse its account of phenomenal qualities. I address these considerations in turn in the next sections.

5.3 - A Challenge to the Policy of Containment

The dual visual systems account offers one lens through which to understand O'Regan's (2011) claim that conscious perception involves a second-order tracking of the regularities manifest in the perceiver's sensorimotor coupling. We can plausibly suppose that the second-order tracking O'Regan describes is carried out by the ventral stream, which, as Clark contends, fails to play a consistently fluid role in the perceiver's ongoing sensorimotor engagements. The object of the tracking is the bodily coupling that comprises the lowest of O'Regan's hierarchically-organised tiers, and this is plausibly modulated by the dorsal stream. Although Noë (2004) fails to say much about perception at the descriptive level employed by

cognitive science, the conscious sensorimotor understanding featured in his account could itself conceivably be the outcome of the same second-order tracking. The Policy of Containment endorses this picture, but proposes that phenomenal qualities are accounted for by the sensorimotor coupling rather than the second-order tracking.

Notwithstanding its explanatory virtues, this account of phenomenal qualities may fail to persuade if it is not tightly integrated into the sensorimotor account of visual processing. In the next two sections, I examine different versions of the same problem and suggest ways the sensorimotor approach can respond. The drawbacks all follow from the apparent role played by the second-order tracking. As I observed in chapters 3 and 4, the variety of sensorimotor mastery deployed in conscious perception - either Sensorimotor Tracking or Sensorimotor Understanding - is not fully developed by the original statements of the sensorimotor theory. The authors explain what the mastery is mastery *of*, namely SMCs, and claim that it is practical and not propositional (O&N; Noë, 2004). Noë (2004; 2012) also discusses its phenomenology at length. However, neither Noë nor O'Regan offer a clear proposal about how this sensorimotor mastery is realised at the descriptive levels employed by cognitive science or neuroscience. The absence of an explicit proposal about how sensorimotor mastery is physically realised is problematic. Firstly, it entails, straightforwardly, an explanatory limitation: the sensorimotor approach, on the present understanding, comes with no clear naturalistic account of what it is for perceivers to be intentionally-directed toward the objects of conscious perception. Moreover, it invites the view that the sensorimotor account must appeal to internal representation, a reading endorsed by commentators such as Roberts (2010) and Hutto and Myin (2012). From this, all sorts of question follow about the nature of the representations and the explanatory work that they properly do.

5.3.1 - The WAY problem

The suggestion that perception involves representation invites a broad threat to the Policy of Containment, namely the charge that perceptual consciousness, in particular the phenomenal qualities that characterise it, must be accounted for not by the law-like patterns of dependence

that govern the perceiver's sensorimotor engagements, but by properties of the representations themselves. In dialogue with O'Regan, Block makes the following point:

In having a cognitive appreciation of a law involving inputs and outputs, one has to think of or represent those inputs and outputs in some WAY. A machine or a creature from outer space might be able to think of human inputs and outputs in WAYS that do not involve any conscious experience. Alternatively, the WAY might itself be phenomenal—say if our cognitive appreciation is coded in imagery. Given that cognition cannot grasp anything without grasping it in some WAY, the appeal to COGNIZING in explaining sensory qualities smuggles in the very notion that is supposed to be explained. (O'Regan and Block, 2012)

This is a fairly broad point, from which we might identify various specific problems. I will discuss two of them in the remainder of this passage, and a third in the section that follows.

One of Block's complaints is that there are ways of representing sensorimotor contingencies that do not involve consciousness, meaning that an appeal to nothing more than the representation of sensorimotor contingencies cannot be sufficient to account for the existence of consciousness. This is a limitation that the sensorimotor theory can and should accept, given that certain other conditions must be in place for an agent to be conscious at all. O'Regan (2011) asserts that consciousness requires the functional realisation of a faculty described as the *self*, while Noë (2004) expresses sympathy toward the view that consciousness is connected somehow to life; and there may be distinct proposals about consciousness that are equally compatible with the sensorimotor approach. While offering some such background conditions for perceptual awareness is no small task, it is not obvious that an account of perception and the phenomenal character of perceptual consciousness is committed to offering a fully-developed answer to this broader question. With the assumption that appropriate background conditions are in place, the sensorimotor theorist can claim not just that sensorimotor contingencies explain differences and similarities in phenomenal quality, but, more fundamentally, that it is by virtue of representations of sensorimotor contingencies that there are, in general, such things as phenomenal qualities at all.

More seriously, Block's objection suggests that once you have accounted for the presence of conscious qualities by appeal to the representation of some feature, the

explanatory weight is borne by some property of the representation other than what it represents - making any proposal about the particular properties represented a dispensable feature of the account. To assume this suggestion is correct without further argument would be a prejudice. Developing the criticism, if it is not the representation's extension that does the explanatory work, it might be an intrinsic property of the neural vehicle that could inhere even if that vehicle lacked other properties, in particular extrinsic ones, by virtue of which it bears representational content.

Block pursues this line of argument elsewhere in an argument against representationalism about phenomenal qualities. In a thought experiment similar to Davidson's Swampman (1987), Block (1998) imagines that an atom-for-atom duplicate of himself appears, by pure chance, which lacks the causal history needed to possess content-bearing inner states. The duplicate nonetheless undergoes the same phenomenal states, suggesting it cannot be content that yields the phenomenal state. Transposing this objection to the sensorimotor theory, Block imagines a neural duplicate of himself that lacks the relevant ongoing or historical sensorimotor interactions but enjoys identical phenomenal states (O'Regan and Block, 2012). Assuming that this is possible, the scenario appears to suggest that ongoing or historical sensorimotor interactions themselves cannot explain phenomenal qualities, because the neural duplicate does not have them. Similarly, the neural duplicate's brain cannot track or represent the appropriate sensorimotor laws, again because there are no sensorimotor interactions to track or represent.

The sensorimotor theory, on a dynamical systems approach, could respond that duplicating the necessary neural states depends on duplicating the appropriate sensorimotor interactions, as we saw in chapter 3. However, the present aim is to defend a more conservative version of the sensorimotor theory that cannot make use of this argument. Another possible response is to endorse disjunctivism, the claim that veridical perceptual experiences are a different kind of state to non-veridical experiences like hallucinations, even if they are phenomenologically indistinguishable. The disjunctivist response would claim that since the duplicates' states differ in kind, with one undergoing a perceptual state and the other something else, we are entitled to provide a distinct explanation for each. This response may fail to convince, however, on the ground that the sensorimotor account of phenomenal

qualities is not meant to explain perception itself, but the phenomenal character common to both the perceptual and the non-perceptual states.

It is perhaps for this reason that O'Regan endorses a different kind of response. He replies that the neural duplicates share a pertinent feature, but, contra Block, that this is not some set of intrinsic neural properties. The pertinent feature is, instead, the fact that the duplicates' brains are each in the state they would ordinarily be in *if* they possessed the relevant ongoing or historical sensorimotor interactions (O'Regan and Block, 2012). This line of response is insufficient to defend the radical claim that the vehicle of perceptual consciousness extends into the outside environment - notably, it is less robust than the argument offered by Hurley (1998) for ECM, which I discussed in section 3.6. But O'Regan's argument is itself sufficient to defend the claim that a proper explanation of conscious qualities must make reference to features of the outside environment.

5.3.2 - Sense and Sensorimotor Contingencies

On a more interesting reading of the WAY objection, phenomenal qualities are not explained by the extension of the representations, but their intensional content, which the empirical work on the ventral stream perception perhaps suggests is likely not to represent the relevant properties *as* laws of sensorimotor contingency. In this section, with special attention to Tye (2000), I will respond to this worry.

The clearest way to explain the distinction between intension and extension is by reference to the appearance it makes in theories of linguistic meaning. In the standard Fregean example, the terms 'morning star' and 'evening star' share an extension, because they both pick out Venus. However, when we first learn that they pick out the same thing, we learn something new. This is explained by the fact that they differ in intension, an attribute of meaning sometimes described as a mode of presentation. A term's intension is often understood to be a set of truth conditions, although the intension could also be some other feature that meets the requirement of explaining how statements of identity between co-referential terms can be informative. The distinction between intension and extension can also be applied to non-linguistic forms of representation, including internal representation at

the explanatory level employed by cognitive science.

An opponent of the sensorimotor theory might concede that for every difference in phenomenal quality there is a corresponding difference in the pattern of sensorimotor dependence governing your bodily interactions with the object perceived. In so far as this holds true, they might say, phenomenal qualities can be individuated by reference to sensorimotor laws. Nonetheless, the objection continues, it is merely incidental that differences in look are accompanied by differences in sensorimotor dependence. The feature that truly explains a tomato's looking red rather than blue, for instance, is not this unextraordinary feature of the extension, but the intension under which the tomato is represented. The intension may not specify that you stand in a particular relation of sensorimotor dependence to the tomato, but instead that you are looking at an object with a particular set of intrinsic properties, or that you are looking at some object which falls under the descriptor 'red'. To secure an absolute victory over the sensorimotor account, an endorser of this line of response must show that the ventral stream, for instance, represents objects under intensions that determinately do not specify patterns of sensorimotor dependence.

The force of this objection will depend on further technical details about the representational content featured in our account of visual consciousness. As it happens, there is at least one leading representationalist account of phenomenal quality pre-existing the sensorimotor account under which intension poses no threat to sensorimotor theory. Although the account offered by Tye (2000) is superficially quite different from the sensorimotor theory, the variety of representational content he endorses makes a good fit, as it will emerge, with the account of phenomenal qualities offered by the particular reading of sensorimotor theory I am pursuing for the time being. In the remainder of this section, I will focus on Tye's account in order to explain how it defuses the threat posed by intension to the sensorimotor account.

The accounts offered by Clark (2008) and Tye resemble each other inasmuch as they both endorse an important role for internal representation in accounting for cognition and phenomenal character. Tye, however, is not directly interested in the details of human cognition, but in using representational content to account for the latter. Where Clark argues that representations enable the exercise of epistemic skills, which themselves account for

phenomenal character, Tye accounts for phenomenal character directly by appeal to the content of the representations. In the apparent converse of Clark's position, Tye ascribes an enabling role to epistemic skills, suggesting that to feature in perceptual consciousness, the representational content must be poised to impact on beliefs and desires. These positions appear to come apart quite considerably, although I will argue later that Clark's account lapses into something close to Tye's than appearances at first suggest.

Tye's account of representational content is in some ways straightforward. His account of content appeals to no more than a simple co-variance relation between aspects of the brain and environment. He suggests that a neural structure bears content about some feature of the environment just in case it is ordinarily present when and only when that feature of the environment is present. Content in Tye's account is individuated externalistically. To understand this point, recall from earlier Putnam's (1975) classic distinction between internalism and externalism about content. The former denotes that content is fixed entirely by features inside the subject's skin. The latter denotes that it is fixed by properties of the extension. To identify which account is correct, we use a thought experiment featuring physically identical inhabitants of superficially identical planets, Earth and Twin Earth. On Earth, the clear liquid that comes up of taps, fills lakes, and so forth, is H²O. Twin Earth is identical in every respect, except the substance bearing the same superficial features is a different chemical, XYZ. If physical duplicates living respectively on Earth and Twin Earth mean different things when they say 'water', then externalism is true; if they mean the same thing, then internalism is true. And the same distinction can be applied to subpersonal content. On Tye's account, then, a neural state represents water if it is ordinarily present when and only when water is present, and it represents H²O rather than XYZ on account of the properties that water happens to actually have.

Suppose that you see a red object, and your brain represents it as falling under the designator 'red'. Further suppose that being confronted by a red object is the same as standing in a bodily relation to the environment where particular patterns of sensorimotor dependence obtain. If externalism is true, then it may follow that the representational content includes this pattern of sensorimotor dependence. However, this suggestion will not by itself serve well enough as a defence of the sensorimotor theory. Even if intension is determined by properties

of the extension, externalism in itself does not require that the intension is determined by every property of the extension. Tye is motivated to endorse intensional content in the first place in part by the need to account for the fact that we may be aware in perceptual experience of some properties of the extension, but not all. In his example, we might see a vegetable as purple, but not see it as poisonous, even in a possible world where all purple things are poisonous. Intension is needed to account for this, since if all purple things were poisonous, the expressions 'purple object' and 'poisonous object' would be coextensive.

To see how the sensorimotor account can be defended, we need to look further at Tye's account. Tye endorses a distinction between *phenomenal* and *epistemic* uses of the word 'looks'. He argues that the distinction is in evidence in everyday discourse about perceptual states when we talk, at times, about an object 'looking F to S' and, at other times, about an object 'looking as if it is an F to S'. The former locution describes the phenomenal character of a person's perceptual awareness, while the latter describes the person's epistemic state. On this account, to say something like 'the shape looks as if it is a trapezoid to Jim' implies that Jim has the concept TRAPEZOID, and believes the object he is looking at falls under that concept. Tye implies that this is neither necessary nor sufficient for us to truly say 'the shape looks trapezoid to Jim', a locution which implies, distinctly, that a trapezoid shape is part of the phenomenal character of Jim's perceptual consciousness. Endorsing this distinction allows Tye to suggest that when we talk about looks we are at least sometimes talking about phenomenal character and *not* epistemic situation. On Tye's account, the shape may look, phenomenally, trapezoid to Jim even if Jim does not possess the concept TRAPEZOID, because the phenomenal look is accounted for by the *non-conceptual* content of his perceptual representations. Tye's motivation for focusing on what we *say*, or are entitled to say, about the way things look to some third person is, presumably, intended as a way to (1) make it plausible that there is such a thing as a phenomenal look which is not to be glossed in terms of conceptual content and (2) discuss the alleged non-conceptual content of perceptual consciousness without attributing the concepts in question to the subject.

Tye describes epistemic 'looks' ascriptions as requiring *hyperintensionality*. Where regular intension allows coextensional expressions to have different contents, hyperintension allows *cointensional* expressions to have different contents. Hyperintension is needed on the

ground that two terms have the same intension so long as they have the same extension by metaphysical or logical necessity. Thus 'H₂O' and 'water' share an intension, as do '55²' and '3025'. These cases are different to the 'morning star' and 'evening star' case, since the shared extensions must be shared in every possible world, while there is a possible world in which 'morning star' and 'evening star' have the intensions they do on this world but are not coextensional, entailing that you are finding out something particular about our world when you learn that they share an identity. The same applies to 'purple object' and 'poisonous object' in Tye's possible world. Nothing turns on whether or not cases where hyperintension is allegedly required cannot, in fact, more accurately be described as instances requiring a variety of plain old intension. It will suffice if we agree that the supposed hyperintension-involving cases are, at least, a distinctive subset of cases involving intension, and I will therefore describe them as involving 'hyperintension' to designate that they belong at a minimum to this special subset, and use 'intension' to describe cases that do not fall under this subset.

Tye argues that the representational content that accounts for phenomenal character is nonconceptual. This reflects the distinction about what we mean in epistemic and phenomenal 'looks' contexts. Crediting Dretske for the example, Tye offers a case in which you see a man in a police uniform, but fail to recognise that he is a policeman. Nothing is strange, Tye observes, about the contention that the man could nonetheless be truly said to look phenomenally to you the way a policeman looks. This lends apparent support to the claim that epistemic and phenomenal looks attributions come apart, and suggests that in the phenomenal case, you need not possess the concept POLICEMAN. If it is indeed appropriate to explain phenomenal look by reference to representational content, then the content in question must therefore be nonconceptual.

Importantly, Tye takes it that the nonconceptual content of perceptual experience cannot be hyperintensional. His argument, which he puts briefly, is that this follows from the fact that hyperintension depends on having "conceptual modes of presentation" (2000, p. 57) which are lacking in the case of nonconceptual content. It is not straightforwardly evident from Tye's discussion why hyperintensional content must be conceptual, but the milder kind of intension that Tye holds is characteristic of nonconceptual content need not be. However,

the view is compelling in light of the following thought, which is perhaps what he has in mind. In the scenario mentioned above, PURPLE OBJECT and POISONOUS OBJECT are co-extensive, because on Earth it just so happens that all purple objects are poisonous, and all poisonous objects are purple. Since the concepts are co-extensive, we cannot identify any physical differences in the way a subject responds to a poisonous object and the way they respond to a purple object. However, we can conceive of a counterfactual circumstance in which the concepts are not co-extensive, and hypothesise that the subject's brain would continue to respond in the relevant way to purple objects that were not poisonous, while failing to respond in the relevant way to poisonous objects that were not purple. The intension is then fixed by the aspect of the extension that the subject's brain reliably responds to, not just now, but in other physically possible circumstances. We cannot, however, make the same move for cases in which the identity is logically or metaphysically necessary, as with water and H₂O, since it is not coherent to hypothesise about how a subject would respond to H₂O that was not water, a scenario that is impossible and perhaps even inconceivable. In the absence of any other readily apparent options for locating hyperintensional nonconceptual content, the only means available to master hyperintensional 'H₂O' and 'water' modes of presentation appears to be by using concepts.

Tye notes therefore that on his account, an object's looking phenomenally red is identical to its looking like an object characterised by the physical properties which constitute being red; as Tye supposes, looking "disposed to reflect such-and-such percentages of light of so-and-so wavelengths" (2000, p. 55). This makes Tye's representationalism quite compatible with the sensorimotor theory's present gloss. On this view, we can take it that second-order tracking of sensorimotor laws involves the activation of representations of the sort described by Tye. In supplement, we merely add that the physical properties represented are not dispositions to reflect such-and-such percentages of light, but patterns of sensorimotor contingency. Importantly, it makes no difference that the sensorimotor laws represented may be more coarse-grained than those discriminated in dorsal stream vision. The sensorimotor account of phenomenal qualities has several important virtues in comparison to Tye's suggestion about the physical properties that underlie colour. Being 'disposed to reflect such-and-such percentages of light' implies that the object activates this disposition when

certain conditions external to the object itself obtain. As such, the disposition is a relational rather than intrinsic property. But Tye's suggestion, which he would of course point out is not intended as a developed proposal, gives no indication of what the other relata is. Sensorimotor theory does offer a proposal on this score, suggesting that the relevant relations are those in which the object stands with ambient lighting conditions and with the agent's body. As a consequence, sensorimotor theory is able to give a better-developed and therefore more principled account of why particular physical states are associated with particular phenomenal qualities, including, crucially, the degrees of similarity and difference between phenomenal qualities. On this basis, there is good reason for an endorser of Tye-style representationalism to adopt the sensorimotor account of phenomenal qualities.

On the present view, the sensorimotor theory emerges as a hybrid of skill-based and representationalist approaches to qualia. The above discussion has shown we can claim, without contradiction, that sensorimotor skill does the main explanatory work in accounting for conscious feel, even though representation is a necessary condition, along with other enabling conditions, for being in states characterised by the relevant phenomenal qualities. On this proposal, the ventral stream plausibly groups together fine-grained patterns of sensorimotor dependence under representations of coarser-grained categories and types, which it deploys for the purposes both Clark and O'Regan recommend. Since the representations lack the strong kind of intension required to represent the relevant features, determinately, *as* categories and types and *not* patterns of sensorimotor dependence, the option remains to individuate the associated phenomenal qualities by reference to sensorimotor law. I have shown in brief that the sensorimotor account, on the present understanding, brings explanatory benefits as an addendum to Tye's representationalism, and another such benefit will become evident in the next section when I discuss the Inverted Earth thought experiment. Returning to the central focus of this chapter, it remains necessary to show that the sensorimotor account is not just coherent but actively compelling when considered alongside the alternative, non-sensorimotor, skill-based approach.

5.4 - Mary and Inverted Earth: A Sensorimotor Response

The above discussion has illustrated that we could quite consistently hold that perceptual experience is realised by the deployment of representations which themselves are not intimately involved with bodily movement, while also holding that the phenomenal character of the experience so-realised is best explained by the sensorimotor laws that govern the perceiver's bodily interactions. Nonetheless, it may seem more natural to suppose that phenomenal character is accounted for by the personal-level skills the representations enable - as Clark (2008) argues, the sorting and sifting of categories and types for the purpose of judgement and action-planning. In this section I offer independent reasons for endorsing the sensorimotor account of phenomenal qualities. Importantly, there is a range of conceptual demands an account of phenomenal qualities must meet that turn on considerations superficially distinct from the character of the representations we are supposing enable perceptual consciousness; and it will become evident as I proceed that sensorimotor and non-sensorimotor skill-based accounts are not equivalent in their ability to meet these conceptual demands. Attention to some classic thought experiments will help show that there is reason to think the sensorimotor approach to phenomenal qualities is more effective at meeting these demands than its competitors.

Consider first the debate over the 'inverted spectrum' thought experiment. Generically, inverted spectra thought experiments feature two perceivers, one whose colour phenomenology is inverted relative to the other, but who exhibits an identical repertoire of behaviour. Block (1990b) has used a version of the thought experiment in an attempt to refute the view, endorsed by Tye (2000) and many others (e.g., Lycan, 1987; Dretske, 1995), that phenomenal qualities reduce to representational content. Block and his opponents share the background assumptions that phenomenal qualities are physical, and that there is no deep-seated conceptual obstacle to accounting for representational content in physicalistic language either. Block imagines a planet which is identical in all respects to our own apart from the fact that the colours of objects are complementary to the colour they possess on Earth, meaning the sky is orange, the grass red, and so forth. The inhabitants of Inverted Earth correctly describe the sky as 'blue', and the grass as 'green', since the meanings of those words

on Inverted Earth differ appropriately from their meanings on Earth. Block asks us to imagine that we put on a pair of colour-inverting glasses and travel to Inverted Earth. The inverting effect of the glasses means that things appear phenomenally to us on Inverted Earth exactly as they do on Earth without the glasses. Yet, at least on externalist accounts of representational content, the content our experiences possess must now differ in line with the differing properties of the objects we are now in causal contact with. This appears to show that phenomenal qualities cannot be adequately accounted for by representational content, at least so long as the content is individuated externalistically.

The sensorimotor approach (on the present gloss) agrees with the representationalists that phenomenal qualities are necessarily individuated by properties represented when we perceive. However, sensorimotor theory avoids the threat that the Inverted Earth case poses to this view, since it is not committed to the suggestion that the intrinsic properties of the objects we come into perceptual contact with could not vary while phenomenal character stays the same. The approach instead holds that phenomenal character will stay the same between Earth and Inverted Earth so long as the relevant patterns of dependence between motor output and sense input are preserved. If this is correct, the inverting goggles could only do the job they are supposed to do if they somehow preserve the patterns of output-input dependence that obtain on Earth. In this event, the phenomenal character of your experience on Inverted Earth is explained, amongst other things, by properties of the objects perceived and properties of the goggles (although, as Myin, 2001, points out, inverting goggles that preserve the relevant patterns of sensorimotor dependence may be a physical impossibility).

Suppose that objects with intrinsic property *X* are, in every possible world, co-extensional with objects with which a perceiver embodied in a particular way stands in sensorimotor relation *S*. This means that a neural representation, lacking hyperintension, does not represent only *X* or only *S*, but represents both simultaneously. On Inverted Earth, the equivalent object has intrinsic property *Y* instead of *X*. Because you are wearing appropriate goggles, neural implants, etc., you stand in sensorimotor relation *S*₂ to the object. *S*₂ is, in one respect, non-identical to *S*, because it is realised in part by an object with different intrinsic properties. However, *S* and *S*₂ correspond isomorphically in a way that *X* and *Y* do not, entailing that there is a level of description at which *S* and *S*₂ are identical. On Inverted Earth,

there is no way to specify determinately whether the neural representation bears content about Y or S2. However, by appealing to the resemblance between S and S2, we can explain why the phenomenology is preserved in spite of the difference in representational content. This response gives a representationalist like Tye further reason to incorporate the insights provided by the sensorimotor account.

The non-sensorimotor skill-based approach offers a different kind of response to the Inverted Earth case. The approach has no need in the first place to insist that phenomenal qualities are individuated by properties of the objects we perceive. All that matters is that, for whatever reason - whether because of the goggles, or some property of the environment - we are able to make a suitable range of discriminations, which do not merely involve judgements that one property is nonidentical to another, but graded judgements about similarity and difference along various scales, for example hue, saturation and brightness. So far, both skill-based accounts offer in principle an equal improvement on representationalism about qualia, although I have established so far that they are not equivalent, since the sensorimotor account appeals, in part, to properties of the outside environment, while the non-sensorimotor skill-based account does not obviously do so.

Consider now a debate concerning the more general question of whether phenomenal qualities are physical at all, and the vivid challenge posed to physicalism by Jackson's Mary's room thought experiment. To recap, the story features Mary, a vision researcher, who is locked in an entirely black and white room and never comes into direct visual contact with any coloured objects. Via black and white TV screens, Mary studies vision science and comes to know every fact that there is to be known about the physical world. Yet when she is finally released from the room, and shown a red rose, Mary, it is alleged, learns a new fact; namely, what it is like to see something red. Since Mary, by hypothesis, already knows every physical fact, Jackson concludes that she learns a non-physical fact, and hence that physicalism about qualia is false (Jackson, 1982). The sensorimotor theory does not come with an entirely prefabricated response to the Mary scenario, but it does invite certain lines of response, and precludes others. The theory says that the look of red can be characterised by reference to a bodily relation the perceiver stands in to the world, which is itself a physical relation, meaning the theory is committed to arguing somehow that phenomenal qualities are physical.

More importantly, the theory is independently committed to the view that to stand in this relation is to possess and deploy a certain kind of skill or practical knowledge. With this in view, it is natural for it to endorse the suggestion (see, Lewis, 1983) that to ‘know what it is like to see red’ is to possess a kind of practical rather than propositional knowledge. Mary’s new knowledge, in this light, plausibly consists of the ability to deploy, at will, the same sensorimotor mastery she deploys to perceive a red object even on occasions where there is no red object present. This is an act of the imagination, and involves naturally extending the sensorimotor theory to encompass imagination in addition to perception. The sensorimotor account is able to make better use of the appeal to practical knowledge to account for the Mary case than accounts appealing to practical knowledge in a more *ad hoc* way, because the theory makes use of practical knowledge to account for perception anyway.

This response gives the sensorimotor account an advantage over the non-sensorimotor skill-based account. In line with Clark’s proposal, that account is committed to the suggestion that seeing something phenomenally as red consists in nothing more or less than having the ability to judge that it is red, or plan one’s actions with sensitivity to the fact that it is red. Mary, locked in the room, possesses all the propositional knowledge about physical facts there is to be possessed. If deploying this propositional knowledge in thought depends on possessing practical knowledge, as Ryle (1949) contended, there is no reason to think she lacks this ability. With this in mind, she can discriminate in judgement the same range of properties before leaving the room as she can after leaving it. If the epistemic skill theorist thinks it is the capacity to make propositionally-structured judgements that makes the difference, they are committed to the view that Mary knows what it is like to see red before leaving the room, a position endorsed by Dennett (1991), but which would, at best, add a substantial and perhaps unpalatable additional commitment to the epistemic skill-based approach.

The non-sensorimotor skill theorist could instead suggest that the capacities involved in ventral stream perception involve neither the deployment of propositional knowledge nor the exercise of fine-grained sensorimotor skill, but a different kind of practical skill. Ward (2009), for instance, suggests that the skill in question is the planning of agential action. If this amounts only to an appeal to coarse-grained rather than fine-grained sensorimotor skill, this is

the position I argued earlier is not, as Clark suggests, incompatible with the sensorimotor theory. Alternatively, the skill might be a capacity for something other than bodily movement or conceptual thought. It is not self-evident what we should take this capacity to be. The most obvious possibility is that it is a capacity for nonconceptual thought. In line with the understanding of nonconceptual content endorsed, for instance, by Tye (2000), a thought of this kind is a conscious state featuring content only characterisable by concepts that the subject does not possess.

However, the appeal to personal-level nonconceptual content is baseless unless some independent considerations make it plausible and intelligible that we actually undergo states of this kind. We cannot make the existence of such states intelligible by appeal to visual phenomenology, as this would be circular: we are trying to make the visual phenomenology intelligible by appeal to the content. The only remaining option is to appeal to the subpersonal content borne by representations realised by the ventral stream or another part of the brain. To explain what it is for a person to experience one patch of red as darker than another, for instance, the present approach would have to say that the person's ventral stream represents it as darker. Notice that the epistemic skill-based account amounts, on this proposal, to nothing more than representationalism in a different guise. Lacking the modification that sensorimotor theory makes on its present representationalist construal, this threatens a return to the Inverted Earth problem, namely the possibility that the discriminations you make could vary while your experience's phenomenal character remains the same.

5.5 - Relation to Rupert's 'Massively Representational Mind'

On its present gloss, the sensorimotor account endorses something similar to Rupert's (2011) notion of *internal-external content*. Typically, representational stories pitched at a subpersonal explanatory level propose that the brain traffics in representations which stand-in purely for outer features of the subject's environment. This holds true in the account offered by Tye (2000), and could correctly characterise Clark's (2008) claim that the ventral stream represents types, categories and locations. Rupert's relatively novel suggestion is that parts of

the brain frequently represent other parts of the brain, and parts of the extra-neural body. In one sense, this content is internalistically individuated, since it is individuated only by features inside the skin of the subject. In a more interesting sense, however, the content is externalistically individuated, because it is fixed entirely by properties of the extension and not by properties internal to the neural state that does the representing. As presently construed, the sensorimotor account suggests something similar. The contents borne by representations in the ventral stream are fixed entirely by properties of the extension. The extension, in this case, is not an external feature but a pattern of dependence governing the perceiver's bodily interaction with the environment. This is not, as with Rupert's internal-externalist content, a purely inner state, but is not wholly external either. Instead it is a relational property that holds between the perceiver and their environment. As I explained in the last section, one benefit of adopting this claim is its ability to overcome the Inverted Earth problem.

It is worth looking in more detail at Rupert's model, the 'massively representational mind', and considering its implications for the sensorimotor account. Rupert offers a very austere account of representation and content, suggesting that X represents Y in the event that it reliably co-occurs with Y. X here is some neural state, and Y is a feature of the outside environment, body or brain. Where the thing represented is another part of the brain, it may be the neural state itself, or it may be the content carried by that neural state. Content is in every case individuated solely by the properties of its extension. Moreover, bearing content is sufficient on Rupert's model for constituting a representation. Rupert makes no demand that the contents fulfil, for example, any particular functional role. In offering such a liberal account of 'representation', Rupert finds that there are representations everywhere; the brain is massively representational. One important feature of the model is that the brain employs a great deal of 'redundant' content, which is content borne at once by multiple representations. He addresses the dual visual systems to illustrate this feature, arguing that the systems carry some content in common, albeit using different representational formats and in the service of performing distinct functions. Rupert does not claim that the ventral stream bears internal-externalist content; he merely suggests that the ventral and dorsal streams each represent the same features of the subject's external environment. However, it is worth considering the dual visual systems in light of Rupert's internal-externalist content and some

of Clark's claims.

One observation Clark makes is that the ventral stream is indeed sensitive to efferent copy from motor output regions in addition to sense input, suggesting it is sensitive to the relation *between* motor output and sense input, i.e. the patterns of dependence appealed to by the sensorimotor account. His argument, however, is that the ventral stream is sensitive to motor output states only during learning phases, when it is acquiring new information about the significance of various kinds of sense input. In the normal course of perception, however, he claims that the ventral stream is sensitive only to sense inputs. The representations it constructs in the course of ordinary perception are, meanwhile, insensitive to *fine-grained* patterns of sensorimotor dependence. Clark's opposition here depends on what he describes as "the best *functional and architectural explanation*" (2008, p. 192) for the available evidence. This is important. The fact that the ventral stream is insensitive to fine-grained differences in sensorimotor law does not mean that it is insensitive to sensorimotor laws. Moreover, the distinction between learning and ongoing sensitivity is only important if we focus on neural architecture. Focusing only on content, where we take this to be explained by co-variance, there is no theoretical difference between a case where the ventral stream co-varies with particular sensorimotor laws because it is responding in part to efferent copy from the motor areas, and a case where it co-varies with sensorimotor laws even though it is responding directly to sense inputs alone.

We could, on Rupert's model, propose that the ventral stream groups the fine-grained patterns of sensorimotor dependence governing dorsal stream activity into coarser-grained representations corresponding to types, categories and relative locations - thereby altering the representational format - and puts them to use for reasoning and action planning. Despite this difference in format and function, the ventral stream represents or borrows content from the dorsal stream. This proposal therefore assigns to the dorsal stream an explanatory role in individuating the content borne by ventral stream representations, even if it is ventral stream representations that serve as the primary vehicles of visual consciousness. If the content by itself explains phenomenal character, then activity modulated by the dorsal stream plays a relevant role in accounting for phenomenal character. And as I argued in the previous section, Clark's preferred variety of skill-based account in the end depends crucially on the content

borne by the ventral stream representations to explain which properties have been skillfully discriminated.

Simultaneous to co-varying systematically with the operations of the dorsal stream, the ventral stream also co-varies with categories and types present in the outside environment. This may lead an endorser of Clark's view to repeat the claim that the ventral stream's co-varying with features of the dorsal stream is merely a contingent feature, while its directly co-varying with features of the subject's outside environment is the pertinent co-variance relation when specifying the content borne by ventral stream representations, and so accounting for the acts of discriminations that explain phenomenal character. In this event, sensorimotor contingencies play no essential role. But nothing in co-variance based accounts of content, like Rupert's, shows that the content of ventral stream representations are not sensorimotor contingencies. The claim that sensorimotor contingencies are the contents is certainly not arbitrary or unmotivated, either, considering that, as Clark appears to allow, the dorsal stream is appropriately described as enabling the perceiver to exercise skillful mastery of sensorimotor laws, and further that the ventral stream, at least during learning phases, does not just happen to co-vary with these patterns but is directly responsive to both to sense input and efferent copy.

It is therefore possible, at least, to characterise ventral stream content as patterns of sensorimotor contingency. There is at least some positive reason to do so in light of the need to offer a coherent account of phenomenal qualities, which may appeal to different features than any accompanying account of cognition. The ability of the sensorimotor model to respond adequately to the Inverted Earth problem makes the sensorimotor perspective in one respect more coherent. On Inverted Earth, the ventral stream begins to reliably co-vary with, therefore represent, different features of the external environment, and therefore enables a different set of epistemic discriminations, where these discriminations are held to involve discriminating types and categories in the external environment. Yet phenomenal character remains the same. However, the ventral stream continues to co-vary in the same way with dorsal stream activity. This makes dorsal stream activity a better candidate for the content borne by the ventral stream than the outside environment, at least when it comes to explaining phenomenal qualities.

5.6 - Chapter Conclusion

This chapter has observed that phenomenal qualities and cognition are separate topics, comprising different sets of question and having the potential to demand distinct, or at least partly distinct, lines of explanation. I have argued that the sensorimotor account of phenomenal quality is not obviously opposed to a representationalist story about cognition, and that the account comes with certain explanatory benefits when contrasted with other available accounts of qualia. I further argued that the epistemic skill-based account of phenomenal qualities, endorsed by Clark, collapses either into an appeal to sensorimotor skill or into a more or less traditional representationalist story, albeit one lacking the explanatory advantage that comes from claiming the properties so represented are patterns of sensorimotor dependence.

Providing a watertight defence, insofar as this is possible, would require surveying further accounts of representational content, in particular internalist accounts, and exploring in more detail ways they can respond to issues arising in debates over phenomenal qualities, in particular Block's Inverted Earth case, which has been a key feature of my discussion in this chapter. Within the scope of this chapter, I have shown, modestly, that the sensorimotor account is not obviously misguided and has some advantages over Clark's account as an account of phenomenal qualities. In the next section, I offer a more robust defence. I will continue to endorse the suggestion that the ventral stream plausibly tracks and perhaps represents lower-level patterns of sensorimotor dependence. However, I will suggest that conscious qualities are accounted for not by the purported contents carried by processes in the ventral stream, but by the coarser-grained sensorimotor skills that this tracking enables.

6 - Sorting, Sifting and Other Sensorimotor Skills

This chapter offers a more radical response on behalf of the sensorimotor theory to the objection levelled by Clark (2008). In this chapter I will suggest that ventral stream representations, if they exist, and whether or not they stand-in themselves for patterns of sensorimotor dependence, enable perceptual consciousness by enabling the perceiver to exercise varieties of sensorimotor skill, which subsume the epistemic skills appealed to by Clark. It is the exercise of this sensorimotor skill that constitutes perception and accounts for phenomenal character. The force of the defence will be based on the sensorimotor theory's ability to account for the intentional directedness and phenomenal character of visual consciousness. This will involve subsuming a range of different skills under the heading of sensorimotor skill, but rather than being an unwarranted instance of, as Clark puts it, the "enactive framework wagging the empirical dog" (p. 192), I will suggest it is a soundly motivated demand that subpersonal theories not mis-describe the relation between, on the one hand, neural and bodily states, and on the other, the personal-level phenomenal states they realise.

6.1 - Avoiding Zombie Intuitions

To begin, let's consider yet another famous thought experiment. Chalmers (1996) is an opponent of the sensorimotor and all other physicalistic approaches to phenomenal consciousness, since he argues that phenomenal consciousness, in general, cannot be physical. To motivate this claim, he argues that we can conceive of a world physically identical to our own and containing creatures that are physically identical to ourselves but entirely devoid of phenomenal consciousness. He call these creatures zombies. On the ground that we can conceive of such a world, Chalmers claims that such a world is metaphysically possible. And if the relevant physical states and processes could exist without the relevant phenomenal states, then the phenomenal states cannot be identical to the physical states. We could attempt to resist Chalmers' argument from the outset by denying that all conceivable situations are possible, or by denying that zombies are even conceivable. Whichever strategy we pick, it

behooves the physicalist so far as possible to reply not just with negative claims but also with a positive proposal about the sorts of physical states that do constitute phenomenal consciousness. Clark (2008) is explicit that he approves of skill-based accounts because they offer a conception of qualia in which we cannot conceive of a possible world of the sort Chalmers describes: on a skill-based account, there is simply nothing more to undergoing a particular phenomenal state than exercising the relevant skill.

Accounts of cognition, some of which comprise explanations of phenomenal character, frequently advocate that certain relationships hold between entities at different explanatory levels. For instance, they might say that representational content, at a subpersonal explanatory level, supervenes on particular kinds of neural state. They may also, like Clark, say that visual consciousness supervenes on the deployment of particular kinds of representation. A supervenience relation implies that the thing doing the supervening cannot vary without the supervenience base varying. At times, a supervenience relation might be metaphysically necessary, meaning it holds in every possible world, or nomologically necessary, meaning it cannot fail to hold in this world. However, there is no general obligation for a supervenience relation to imply anything so strong. For instance, it would be informative and plausibly true to claim that visual consciousness supervenes on ventral stream activity in presently living humans, on the ground that visual consciousness in humans cannot vary without ventral stream activity varying. However, it would be unjustified to claim, with confidence, that future humans could not in principle use technological artifacts or a different brain areas to realise the same kind of consciousness.

Accounts of phenomenal quality are also entitled, for some explanatory purposes, to make supervenience claims without committing to their obtaining in every possible world. To show that zombies are impossible, however, an account of qualia must offer something more robust. At a minimum, it needs to show that a given physical property, or physically-realizable functional property, is metaphysically sufficient for an appropriate phenomenal state to obtain. Logically speaking, metaphysically sufficient conditions of this type are enough to entail that there is no possible world containing zombies. However, metaphysically sufficient conditions for phenomenal consciousness would not be enough to discount the inverse view that there is a possible world where phenomenal consciousness exists without any

attendant physical state. There is also the question of how to convince present day dualists subscribing to the zombie intuition that a given set of conditions is indeed metaphysically sufficient for phenomenal consciousness, considering we cannot travel to other possible worlds to find out. To offer a convincing reply to dualists, we must make it convincing that the physical states under consideration are not just sufficient for, but *identical to*, the appropriate phenomenal states. This will make it inconceivable that possible worlds exist where the physical states in question are not metaphysically sufficient.

One view is that psycho-physical identities require metaphysically necessary as well as sufficient physical states. To identify metaphysically necessary conditions, we must, of course, be indifferent to facts that are merely contingent. Assuming, for example, that the ventral stream itself only plays a contingent role in realising the cognitive processing that underlies visual consciousness, the metaphysically necessary conditions will not include having a ventral stream. An endorser of a Clark-style position would claim, more plausibly, that physically-realizable functions are what count. Thus, human-style visual consciousness depends by metaphysical necessity on representations carrying out the functional roles performed in this world by the ventral stream. These include playing an appropriate functional role in the wider cognitive economy, and bearing content that is systematically insensitive to occurrences at the sensorimotor periphery. If Clark were indeed arguing that these states are metaphysically necessary for visual consciousness of the sort undergone by human perceivers, he would be making a strikingly strong commitment. On this basis, he would have to claim that a subject undergoing the same functionally-defined inner activity but altogether lacking the capacity to integrate this activity with skillful behaviour will still experience all or most of the same visual states we do. He would also have to claim that a subject lacking the functionally-defined inner states we possess, but reliably manifesting an identical repertoire of skillful bodily behaviour, necessarily lacks the kind of visual consciousness we enjoy. These claims might be correct, but they go significantly beyond claims that can be straightforwardly justified by attending to the workings of the ventral stream.

To make the sensorimotor approach a more compelling alternative, we must be careful how we frame it. The approach would be committing itself to something considerably less palatable than the above proposal if it claimed that bodily activity is sufficient for visual

consciousness even in a metaphysically possible world where the perceiver lacks a brain. It would not be compelling, either, to insist that bodily activity or activity in motor areas of the brain are necessary for visual consciousness in every possible world. What the sensorimotor approach ought to say, instead, is that the metaphysically necessary and sufficient condition for having visual consciousness characterised by the appropriate phenomenal qualities is the deployment of inner states that make an intelligible contribution to enabling appropriate bidirectional patterns of systematic dependence between sense inputs and motor outputs. Nothing in Clark's (2008) critique demonstrates that ventral stream activity could not be made intelligible this way. If successful, the present approach falls short of establishing the strong claim that the extra-neural body and environment play a metaphysically necessary role in realising perceptual consciousness. But, crucially, it offers a way to resist Clark's claim that the skills which constitute visual consciousness are not sensorimotor skills. The present picture has the advantage of allowing that perception *could* be sensorimotor, while the Clark-style position must actively deny that a creature lacking the right internal states could undergo the kind of visual consciousness humans do, even if they possess the appropriate sensorimotor skills

If there is a need to choose whether sensorimotor or non-sensorimotor skills are more plausible candidates as metaphysically necessary conditions for perceptual consciousness, the sensorimotor account thence has an advantage. However, Clark could reject the need to make this choice. Instead, he might claim that we should endorse pluralism about the skills that comprise the metaphysically necessary conditions. On this understanding, the skills we must appeal to in order to dissolve the zombie intuition are neither sensorimotor nor non-sensorimotor, but non-specific capacities to make the appropriate discriminations. The pluralistic perspective entails that the generic skills in question could in principle be constituted, in different cases, by quite different capacities. In this light, Clark's objection to the sensorimotor account can be taken as targeting the claim that the skills in question are constituted by sensorimotor skills in this world. This returns the advantage to Clark.

However, certain considerations blunt the force of this response. It is an open question whether sensorimotor and non-sensorimotor skills belong to one broad kind. The capacity to make skillful bodily movements and the capacity to represent types, categories and

locations are, on the face of it, quite different kinds of capacity. The only thing they necessarily have in common is that they involve responses that co-vary systematically with aspects of the outside environment. But lots of things exhibit this kind of sensitivity and do not justify ascriptions of 'skill'. A skill-based account of qualia must therefore appeal to some further criterion. In Clark's account, this is the content and functional role of the representations deployed by the ventral stream, however these notions are themselves glossed. In the sensorimotor theory, the skill is ascribed by properties of the outward behaviour, a claim I will elaborate on in the next chapter. Once these criteria are added, there is no guarantee that the skills appealed to respectively by sensorimotor and non-sensorimotor skill-based accounts have any common factor that would justify our thinking of them as two different ways the same generic skills could be constituted. In offering psycho-physical identities, it will not do, then, to claim that skills are the metaphysically necessary conditions for perceptual consciousness without specifying in more precise terms what sort of skill we have in mind. As I have argued, sensorimotor skills are a more compelling candidate, since they accommodate the possibility that perceivers lacking our functionally defined inner states may nevertheless enjoy similar types of experience.

Adopting a different tack, we might find that psycho-physical identities fall short of demanding metaphysical necessity. We could instead make a case that there is simply nothing more to undergoing a given phenomenal state than exercising a particular skill on the ground that exercising the relevant epistemic skills is, in this world alone, even in one subject alone, constitutive of undergoing the phenomenal state. The constitution claim is less strong than the metaphysical necessity claim, since it allows that a given experience might alternatively be constituted of quite different things. For example, there might be nothing more to your travelling to work today than your cycling from A to B, but tomorrow there might be nothing more to your travelling to work than your driving from A to B. Constitution claims go beyond mere sufficient conditions, however, by showing that there is nothing more to consciousness than exercising the appropriate skills, in this world at least. In this light, an endorser of the non-sensorimotor, skill-based account might accede to what we could describe as a disjunctive account of the skills that are metaphysically necessary for visual consciousness. This resembles the pluralistic account of skill I considered above, but rather than claiming that sensorimotor

and non-sensorimotor skills belong to a single broader class, it simply claims that the relevant skills could either be sensorimotor skills OR non-sensorimotor epistemic skills. This avoids the need to claim that a possible creature possessing sensorimotor skills like ours but lacking the relevant inner representational states could not possibly share human-style visual experience, but allows Clark to maintain that in human perceivers, in this world, the skills are non-sensorimotor.

This view has more promise. However, it depends on the claim that there is a genuine disjunction in this world between sensorimotor skills and the epistemic skills apparently enabled by the ventral stream. To show that such a disjunction obtains, Clark would have to show that there can be a difference in epistemic skill without a difference in sensorimotor skill, and vice versa. To prove the sensorimotor account wrong, Clark would also have to show that phenomenal qualities vary in line with the former and not the latter. Clark (2008) does not establish this, only demonstrating that not every difference in sensorimotor skill comes with a difference in phenomenology, an observation the sensorimotor account can accommodate.

If the sensorimotor perspective is correct, there is a constitutive link between perceptual consciousness and certain bodily capacities. It might prove that the bodily capacities in question are constituted in turn by processing involving that deployment of representations that stand in for types, categories and relative locations. But it would, in this eventuality, be the bodily capacities themselves that rightly overcome the zombie intuition. In the next section, I will defend the conceptual claim that the epistemic activities described by Clark, the sorting and sifting of types and categories, in fact constitute sensorimotor skills.

6.2 - Sorting and Sifting as Sensorimotor Skills

Clark (2008) argues that if conscious perception, as realised by the ventral stream, involves sorting and sifting of types and categories in place of the fine-grained sensorimotor behaviour enabled by the dorsal stream, then the sensorimotor approach is thrown into significant doubt. This only follows, however, if sorting and sifting (and the other 'epistemic' activities mentioned, like classifying, selecting and choosing) are not varieties of sensorimotor activity themselves. If they are, then it is straightforward that they draw on mastery of the laws of

sensorimotor contingency. Clark does not explicitly address this possibility, but his rejoinder to the sensorimotor approach implies that sorting and sifting draw purely on internal processes, in particular the activation of internal representations, and not sensorimotor mastery. In this section, I am going to attempt a defence of the sensorimotor approach by showing that sorting and sifting are themselves plausibly best understood as varieties of agent-level sensorimotor activity.

Sorting and sifting make an interesting case study, since they are paradigmatically cognitive activities, while at the same time being, ordinarily, activities carried out through bodily engagement with the environment. This point dovetails with the main focus of Clark's (2008) book, which argues that many of the cognitive processes traditionally held by cognitive scientists to depend solely on neural activity also depend on the agent's bodily engagements with the outside environment. Clark claims that in some cases, the substrate of cognition actually extends into the outside environment. In other cases, he merely claims that to correctly understand how a subject performs a given task, we must make explanatory reference to the subject's embeddedness within a specific environmental niche.

A reasonably paradigmatic example of the latter, offered by Clark, concerns a technique employed by a practised bar tender (p. 62). Each drink on the menu is served in its own distinctively shaped glass, and as the orders come in, the bar tender places appropriate empty glasses in a spatial sequence corresponding to the order in which the drinks are to be served. This saves the bar tender the trouble of trying to remember which drinks are to be served and in which order. Clark stops short in this case of claiming that the bar tender's manipulation of external props actually constitutes part of the subject's cognitive processing, but claims that the trick contributes in an important way to the bar tender's performance of the task, while altering the character of the cognitive tasks that must be carried out by the subject's brain. It does not matter too much for my present point whether the bar tender's interaction with their niche is supposed to itself be cognitive. It suffices to notice that the subject, here, is performing tasks of the kind allegedly carried out by ventral stream perception - classifying, selecting, choosing, and so forth - by means of their sensorimotor coupling with the outside environment. It is striking that from a commonsense perspective, this is no surprise, since in everyday language, if not the language of cognitive science, activities like

sifting and sorting are almost always activities carried out through the bodily manipulation of external props.

By extension from the bar tender case, it is worth entertaining the possibility that the sorting and sifting of types, categories and locations that Clark thinks underlies perceptual consciousness are also constituted by bodily interaction with the environment. This proposal is not altogether outré by the standard of the dual systems model. The model describes dorsal stream activity, which is not supposed to come with any attendant consciousness, as a kind of 'vision', reserving the term 'perception' for cases associated with conscious awareness. If the dual systems and sensorimotor theories are agreed that there is an unconscious kind of vision best characterised as a sensorimotor process rather than by an orthodox representational model, then relative to accounts that reserve 'vision' for states or activity realised solely by the deployment of internal representations, it is a lesser step to contend that conscious perceptual awareness is also constituted by the exercise of sensorimotor skill.

With the above points in mind, the sorting and sifting identified with perceptual awareness is plausibly constituted - in *some* cases - by successful behaviour directed at a practical goal. For example, your sorting red from green might involve nothing more than your appropriately pulling your car to a stop in front of a red traffic light. It would be useful, however, to give some account of how such behaviour is actually achieved. Such an explanation can be given in part, by reference to bodily activity directed toward an epistemic goal, for example angling your head upward to see the traffic light. This sort of ongoing re-alignment of bodily posture for the purpose of gathering information also plausibly constitutes part of the process of sorting and sifting. Although perceptual consciousness does not usually involve active manipulation, with your hands, of external props, a core analogy with the bar tender case is maintained by virtue of the fact that the perceiver, in line with O'Regan's idea of the world as an outside memory, makes persistent changes to their bodily relation with the environment in order to achieve both epistemic and practical goals. The fact that bodily movement plays an important role in perception does not, of course, prove that the extra-neural activity actually constitutes, rather than merely assists, the exercise of the epistemic skills Clark talks about. However, my proposal ought to make it at least conceivable that identifying colours, shapes, locations and so forth is best characterised by reference to a

bodily process.

This suggestion should not be confused with the implausible thesis that perceptual consciousness is realised by nothing more than outwardly-observable bodily behaviour. The proposal is that activities such as sorting and sifting are constituted by sensorimotor coupling carried out by the whole organism, drawing on both neural and non-neural activity. Where outward activity is present, this alone is sufficient, sometimes, to constitute an act of sorting or sifting, provided it is enabled in some intelligible way by the subject's inner activity. To account for discriminations made in perceptual consciousness that go beyond those which can be individuated by outward bodily activity, including documented experiments where subjects have had their whole bodies, including eyes, completely paralysed while enjoying a restricted degree of perceptual awareness anyway (Whitham et al., 2011) a different story must be given. However, these cases could be accounted for by appeal to the agent's *preparedness* to make skillful bodily movements.

It is conceivable, of course, that in order to be prepared to act, the subject must deploy internal representations, such as forward models or action-oriented representations like the ones endorsed by Ward, Roberts and Clark (2011). This is not necessarily a problem for the sensorimotor theory, which on some formulations accepts that subpersonal representations play a role. In this event, perception can still be understood constitutively as a temporally-extended process, which, aside from exceptional circumstances, incorporates bodily movement, which forward models help enable. This is compatible, for instance, with O'Regan's (1992) original presentation of the outside memory idea, which suggests that bodily exploration is assisted by sparse representations of objects and their egocentrically-specified location in the subject's environment.

However, recent work on preparatory activity in the motor cortex supports the suggestion that neural representations are not necessarily involved in preparing for behaviour. It is agreed that neural activity takes place in the motor cortex prior to movement which helps enable the subject to move their eyes, limbs (etc.) the appropriate distance, direction and speed. A popular view has it that certain neurons become weakly active in preparation for a movement, and that the movement takes place when that neural activation reaches a certain threshold, suggesting that the preparatory activity acts as a representation of the movement

that is about to follow (e.g., Tanji and Evarts, 1976). However, this assumption is challenged by recent empirical work. One study recorded neural activity in monkeys prior to and during specific reaching tasks. The study found that the preparatory activity correlated with a given movement was quite different to the neural activity that occurred while the same movement took place. This does not establish that there is no representation going on, but it counts against the representational view I just described. Moreover, it is expected by a mechanistic account, in which the preparatory activity does not represent the movements that are about to follow, but initiates a dynamical system in which the body movements takes place (Churchland, Cunningham, Kaufman, Ryu and Shenoy, 2010). Given that the ventral stream has the function of enabling coarse-grained sensorimotor behaviour, it is conceivable that its activity could also be understood as non-representational preparatory activity for embodied acts of sorting and sifting. Moreover, this shows that perceivers could be engaging the relevant bodily capacities even where the capacities are not manifest in outwardly observable bodily movements.

Sorting and sifting are, in general, naturally understood as agent-level sensorimotor activities. In the domain of conscious perception, it is tenable, at least in some cases, that sorting and sifting are constituted by agent-level sensorimotor coupling. This ought to defuse the suggestion that where perceptual processing involves sorting and sifting, it does not involve sensorimotor skill. My discussion has aimed to offer some positive motivation for the view that sorting and sifting in perception is indeed always a sensorimotor process. It remains for the non-sensorimotor skill-based approach to show that sorting, sifting, classifying and selecting can be explained purely by reference to what the brain does without committing a category mistake.

6.3 - Homunculi that Sort and Sift

In the last section, I noted that sorting and sifting are, on an everyday understanding, properly thought of as whole-agent bodily interactions with the outside environment. Clark's (2008) account correctly acknowledges that the skills at play in a skill-based account of qualia are personal-level capacities, meaning they are possessed by the perceiver as a whole, not parts of

the perceiver. It is hard, however, to imagine what the personal-level capacities are if they are not sensorimotor skills. Clark's account is mainly pitched at the subpersonal level, and the only way to make sense of sorting and sifting being wholly inner activities is to suppose that they are carried out by some purely-inner parts of the agent, in particular the agent's brain, and become personal-level activities when the conditions are in place for these subpersonal activities to enter into consciousness. The burden of proof is on the endorser of Clark's non-sensorimotor skill-based approach to show that there are or could be any purely internal activities that can be legitimately described as sorting or sifting.

Drawing on Bennett and Hacker's (2003) view, one might argue that everyday usage is sufficient to show that 'sorting' and 'sifting' can only meaningfully describe behaviour by a whole agent. In fact, 'sorting' and 'sifting' are, in everyday practice, meaningfully applied to artifacts such as coin sorting machines and sieves, hence they fall under the limited category of intentionalistic predicates that Bennett and Hacker (2003) allow can be applied both to persons and other things. However, the Bennett and Hacker-derived point, if I were pushing it, would nonetheless apply to many of the other activities I have been subsuming under the phrase 'sorting and sifting', such as classifying, selecting, choosing and comparing. In fact, for reasons I discussed earlier, the view that everyday linguistic practice offers a definitive guide to the admissibility of applying intentionalistic or other vocabulary to a particular object is contentious, and there are apparently strong reasons to think it is false. My present line of argument does not pursue this view, but depends on the observation that sorting and sifting are as a matter of *fact*, rather than a matter of linguistic practice, commonly carried out by whole agents in interaction with the environment. All the same, any deviation from ordinary linguistic practice requires, as I argued earlier, an accompanying justification for the application of the relevant terms outside of their usual contexts.

To make a compelling case, the non-sensorimotor skill-based account must endorse the suggestion, familiar from Dennett's (1978) proposal concerning homuncular decomposition, that sorting and sifting are practical activities carried out by *subpersons* (i.e. homunculi). If this subpersonal account succeeds, subpersons could engage in their practical tasks without typically having any attendant effect on, or conceptual connection with, the whole agent's bodily interaction with the environment. However, the account must offer

compelling reasons to think these subpersons exist. Consider that when we observe the bar tender, we can tell, unproblematically, that sorting and sifting is taking place. Typically, agent-level, embodied cases like cocktail-making involve ongoing physical manipulation of the objects being sorted and sifted. Crucially, such processes must have a meaningful outcome in the sense that we can, if we like, evaluate the outcome relative to some success conditions, for instance whether or not the bar tender has assembled the correct drinks. Perception is a little different, since the perceiver need not move objects around in order to perceive. However, as with the bar tender case, it should at least be uncontroversial that we are frequently licenced to ascribe visual and/or perceptual states to agents by virtue of their outwardly observed skillful bodily engagements, leaving open the possibility that the outward bodily engagements in question constitute part of perceptual processing.

If we wanted to defend Clark's position, one thing we would have to show is that the subpersons I just mentioned can, in general, be ascribed functions, which might include sorting and sifting. One feature of cocktail-making, and of the kind of skillful bodily behaviour that allows us from the outside to ascribe perceptual states to subjects, is that it is goal-directed. To maintain in a demanding way the link between the posited subpersonal activity and the agent-level cases, the activity by the subpersons must be goal-directed too. To establish this, we would have to show not only that the subpersons happen to sort and sift, but that they have the proper functions of sorting and sifting, which is to say that they have sorting or sifting as their purpose; and the theoretical feasibility of ascribing proper functions to biological things, such as parts of the brain, is a matter of controversy. On this question, one popular line of thought has it that Darwinian natural selection offers a way to naturalise teleology, on the basis that a trait has a particular function just in case its having that function explains the trait's persistence or the organism's biological fitness. The diametrically opposed position is that natural selection - by contrast with divine creation, for example - entails that there is, properly speaking, no function or design in biology at all (Allen, 2009). Even this latter view does not, however, rule out the possibility, (i) that neural systems just so happen to perform the tasks under consideration, even if they are not proper functions, and (ii) that proper functions can be legitimately ascribed as metaphors, if not literally.

Perhaps the function issue is not of crucial importance in deciding whether or not the

ventral stream could sort and sift. For the purpose of establishing that subpersonal systems sort and sift in a way sufficiently analogous to the behaviour of whole agents, we may not actually need to establish goal-directedness in a robust way, since it is not obvious that goal-directedness is a necessary rather than contingent feature of agent-level sorting and sifting. Even if one were sceptical about proper function, then, we could posit that the ventral stream just so happens to sort and sift, with the addendum, if we wanted to endorse it, that sorting and sifting are its proper functions metaphorically speaking.

Let's take it that the aprioristic considerations allow that the idea of ascribing functions like 'sorting' and 'sifting' to neural subpersons is uncontroversial. Even where this is allowed, there persists a distinct obligation to show that there is something purely inner that can be evaluated relative to success conditions. Consider that in the first place, the homuncular subsystems require an input, i.e. something to sort or sift. These could be representations, but they could also be raw sense input signals; this is unproblematic. However, the systems must, more significantly, also have an output, some outcome by virtue of which sorting and sifting can be said to have taken place. The internalist model denies that these outputs are external to the agent as a whole, although to maintain a comparison with 'sorting' and 'sifting' in the ordinary senses of the words, the objects must be external to the homunculus.

The output of an act of sorting or sifting could, in principle, be an alteration to the same state that served as the input, as with the bar tender case, where the cocktail glasses, in one arrangement or another, feature as both the objects and the outcomes of the sorting and sifting. Since perception involves processing an ongoing stream of raw data, the homuncular subsystems that receive this input more plausibly send outputs elsewhere, i.e. to other subsystems. To licence the claim that there is sorting and sifting going on, there must be some principled way of evaluating the outputs, and this entails that the outputs must have meaning. A raw sense input signal is not intrinsically meaningful, but a representation is, by virtue being intentionally directed toward something external. This means that the non-sensorimotor, skill-based account must be accompanied by an account of representation that is both plausible and robust enough to properly account for intentionality.

6.4 - Radical Anti-Representationalism

On a 'radical' reading, the sensorimotor account is situated some way closer to the original enactivist approach that originates with Varela, Thompson and Rosch (1991), and closer still to the dynamic account of Hurley (1998), the radical embodied cognition of Chemero (2009) and the radical enactivism of Hutto and Myin (2012). The most important thing these approaches have in common is that for one reason or another they propose that perception is realised in part by extra-neural activity and takes place without internal representation. If this perspective can be convincingly motivated, either by reference to specific empirical evidence concerning perception or on broader theoretical grounds, then the sensorimotor theory can respond head-on to Clark's (2008) objection that the cognitive processing underlying perceptual awareness does depend on internal representation and not bodily movement.

Enactivism, in its best established guise, originates with Varela, Thompson and Rosch's (1991) book *The Embodied Mind*, drawing also on phenomenology (e.g., Merleau-Ponty, 1945) and earlier work by Maturana and Varela (e.g., 1980) on the biological phenomenon of autopoiesis, the process of active self-creation through which cells and larger organisms maintain their own internal structure. This tradition denies that the world, as it features in the subject's perceptual experience, results from internal representations of an objectively-specified pre-existing environment. It instead claims that the agent creates its own environment and experiential world through the same process of bodily coupling with its surroundings by which it generates and maintains its own existence. Perception, by this right, is inextricably linked to action. Even more distinctively, cognition is on this variety of enactivism inextricably linked with life (see, e.g., Thompson, 2007). The neighbouring accounts I discuss are not committed to this tradition's view that the organismic interactions that constitute perception necessarily overlap with the processes by which the organism constitutes a living thing. For this reason, they are not necessarily committed to the view that perceivers must be alive at all. However, Hurley, Chemero, Hutto and Myin, and the sensorimotor account on a 'radical' gloss, all concur that perceptual consciousness is realised by sensorimotor engagements and not representations.

Hutto and Myin (2012) adopt the 'enactive' label for their own approach, which they

call Radical Embodied or Enactive Cognition (REC). REC is motivated by the claim that 'basic' cognition, meaning any cognition not structured by natural language, 'doesn't need' and 'can't have' representational content. The authors argue that all attempts to account naturalistically for content fail because they are unable to surmount the so-called Hard Problem of Content, a deep-seated conceptual barrier to explaining how one physical state of affairs can bear truth or success conditions relating to another. The root of the problem, the authors claim, is expressed by their headline principle that Covariance Doesn't Constitute Content. This is to say that X co-varying with Y does not entail that X bears content about Y, however we flesh out the precise details of the covariance relation. For instance, a tree's rings reliably co-vary with its age and therefore carry information about its age, where information is understood as a kind of covariance. However, this does not entail that the tree trunk bears content about the tree's age, as content in this case requires truth conditions, and nothing about mere covariance entails truth conditionality. An opponent could bluntly insist, of course, that covariance *does* entail truth conditionality, to which Hutto and Myin would reply that this violates explanatory naturalism by introducing an unexplained and 'spooky' new entity.

The introduction of additional theoretical apparatus, in particular biological function, does nothing to overcome the problem so long as we continue to appeal to a covariance relation to do some of the necessary explanatory work. If Hutto and Myin's 'Covariance Doesn't Constitute Content' principle is correct, this at once undermines a large category of accounts of cognition, among which are developed attempts to account naturalistically for content (e.g., Dretske, 1981) and theories that appeal more or less in passing to the notion that covariance does constitute content (a couple of examples, which I discuss elsewhere in the thesis, include Tye, 2000; Rupert, 2011).

Hutto and Myin broaden their attack on content by targeting leading alternatives to covariance-based accounts, in particular the producer-consumer model proposed by Millikan (1984). Millikan's account suggests that content obtains in informational transactions between a 'producer' and a 'consumer', which might be different creatures or different parts of the brain. This is a variety of homuncular view, which I discussed in the last section and, earlier on, in chapter 2. Content is individuated, at least in large part, by appealing to the

purposeful manner by which the consumer responds to given pieces of information. Purpose can be accounted for naturalistically by reference to the consumer's evolved biological function. For instance, a frog's motor output system might respond to a particular input by causing the frog to dart out its tongue. If this response evolved because it allowed the frog's ancestors to eat food, then the input received by the frog's motor system carries content about food.

Hutto and Myin claim that the reference to a 'consumer' suggests that there is pre-existing content to be consumed, a violation of the principle that Covariance Doesn't Constitute Content. Where biological function does all the explanatory work, they suggest the 'consumer' would be better labelled a content 'creator'. Assuming that Millikan does mean biological function to do all the work necessary in individuating content, Hutto and Myin argue the account fails anyway. The argument, which they attribute to Fodor (1990), notes that content must involve intension as well as extension, since if there were only extension the purported content could not be false, and if it cannot be false then we cannot ascribe to it any specific set of conditions under which it is true. Natural selection, however, is only sensitive to organismic dealings with the outside environment, i.e. the extensions of the purported representations. It therefore provides no resources for explaining intension. Millikan (1993) has responded that the truth-conditions are not specified by current properties of the extension, but by properties of the extension that obtained in the historical environment where the trait evolved. However, Hutto and Myin contend that this does not solve the problem. They say: "Even if we can specify what is meant to be targeted that would give us exactly no reason to think that the targeted item is represented in a truth-conditional, referential, or otherwise semantic way - i.e. that it has intensional content" (2012, p. 80).

Their point, as I take it, is that Millikan's appeal to teleology is sufficient to show that a frog may dart out its tongue in response to a fly, or alternatively to some other object resembling a fly, using a faculty that has the proper function of responding only to flies. However, teleology is not sufficient to show that any part of the frog is evolved to bear intensional content, considering that natural selection could not be sensitive to the guise under which flies are represented, so long as the frog successfully eats them. Without intension, we could legitimately describe a frog's darting out its tongue in response to an

inedible object as an unsuccessful attempt to perform a proper function, namely to eat. But we could not make the further claim that the frog or some part of the frog represents the inedible object as food, because we have no way to specify whether it represents it as food, or as a fly, or whatever.

I take it that Hutto and Myin make a compelling case that no attempt to naturalise content has so far been given that will satisfy their metaphysical demands. One way to pursue the debate is to examine these and other attempts in further detail, and to investigate further attempts that could be made. Fruitful though this may be, a couple of striking facts undercut this. One is the sheer number of contemporary accounts featuring representation and content that appeal to nothing more than reliable covariance (e.g., Tye, 2000; Rupert, 2011). If some different or more baroque attempt to naturalise content eventually proves correct, there is no guarantee that these content-involving approaches will be able to straightforwardly substitute the new account of content for the covariance-based account while leaving all their other key claims intact. With this in mind, Hutto and Myin will have made substantial progress toward making radical enactivist approach to cognition appealing if they can successfully unseat covariance-based accounts of content.

Offsetting this argumentative advantage is the possibility that Hutto and Myin are simply making unduly grand metaphysical demands of their opponents. Although Hutto and Myin probe the specifics of individual attempts to naturalise content before arguing that they fail, their position is not altogether different from Bennett and Hacker's (2003). Hutto and Myin, for their part, assume without much argument that the meaning of words including 'represent', 'content' and 'truth condition' are such as to place a demanding metaphysical burden on anything that can be correctly said to fall under those terms, and go on to examine specific proposals before finding that they are unable to meet this metaphysical demand. Bennett and Hacker, by contrast, assume without much argument that there can be no underlying feature shared by persons and parts of persons that might actually licence the ascription to both of predicates like 'representing'. Where Bennett and Hacker make a linguistic demand, they appeal implicitly to a metaphysical assumption. Hutto and Myin make a metaphysical demand, but appeal implicitly to similar linguistic assumptions. One theorist who disputes this approach is Chemero (2009), whose 'Radical Embodied Cognitive

Science' defends a similar positive programme to Hutto and Myin's, but who sharply disagrees with their negative claims.

Chemero argues that in line with the linguistic conventions employed by scientists, the words 'representation' and 'content' make only modest metaphysical demands. By this latter standard, certain kinds of covariance relation do, in fact, constitute content. Hence, we could well say that a tree's rings represent or bear content about its age. Defending content-free cognition, Chemero more modestly insists that cognitive scientists should only appeal to representation or representational content if doing so provides the best explanation for a given phenomenon. This, he takes it, is an 'epistemic' rather than a 'metaphysical' demand. Hutto and Myin say something superficially similar when they concede that the appeal to representational content would be legitimate if the explanatory success of perceptual science depended on it. However, unlike Chemero, Hutto and Myin suggest that even where the appeal to content is motivated by epistemic rather than metaphysical considerations, it nonetheless entails a robust metaphysical commitment of the sort Chemero rejects.

In one way or another, each side justifies its stance by appeal to naturalism. For the present purpose, we may take it that the naturalism in question is nothing more robust than the minimal naturalism advocated by Wheeler (2005) and endorsed by Hutto and Myin. This amounts to the claim that science and philosophy must be consistent, and that where their respective claims conflict, philosophy must give way. Chemero argues in effect that where scientific accounts find theoretical utility in positing 'representations', and apply the word in a manner consistent with their own linguistic convention, philosophical considerations cannot unseat them, since philosophy must give way to science. In this sense, Chemero is assuming the mantle of explanatory naturalism.

Hutto and Myin deploy minimal naturalism to the opposite effect. They take it that notwithstanding the apparently productive deployment of the terms 'representation' and 'content', especially in empirical work, no present theory shows how the physical states in question genuinely represent or bear content in light of the metaphysical demands Hutto and Myin think apply. Their point, to extrapolate a little, is this. Explanatory utility is insufficient to licence the claim that some entity exists if there is otherwise no empirical or theoretical evidence that the entity does, or even could, exist. If there is no reason to think that an entity

exists, it is unscientific to claim it does exist. On this basis, the authors appear to have naturalism on their side.

In Chemero's favour, it is evident that scientists, especially neuroscientists, frequently use 'representation' and even 'content' to denote nothing more demanding than covariance. Moreover, it is true that people are entitled to use words in accordance with the linguistic conventions in operation within their own discipline. In Hutto and Myin's favour, the problem is that 'content' is also frequently used to denote something metaphysically more robust than mere covariance, even on accounts that appeal to nothing more than a covariance relation to account for it (Tye, 2000, is a good example, since his account of phenomenal qualities depends strongly on the idea that they reduce to representational content and not mere co-variance relations). Where content is used to denote a covariance relation, it is, given this equivocation between speakers, not always apparent whether the speaker means to make an extravagant metaphysical claim about the covariance relation in question, or means to use 'content' in a deflationary way to denote something metaphysically austere.

To resolve this impasse, both sides ought to endorse McDowell's (1994b) position that where content is ascribed subpersonally, i.e. to parts of persons rather than persons, it should be understood by default, and in all cases, as no more than a metaphorical ascription. 'As if' content, as he calls it, is not really content, but may be legitimately ascribed just in case it proves useful or indeed indispensable as a theoretical posit to explain information processing transactions between an agent's parts, a possibility he happily endorses. By adopting this as a permanent disclaimer, it is possible to accommodate Chemero's position that representation and content are legitimate posits when deployed in a metaphysically austere way. At the same time, this ought to resolve any suggestion that 'content', deployed in a superficially harmless way, actually disguises a more robust metaphysical commitment.

Although metaphorical content is admissible in principle, it is subject to certain provisos. One is that if we aim to account for phenomenal qualities naturalistically in a way that defeats anti-physicalist intuitions, we cannot do so by appealing to solely to metaphorical content. Although a physical state that merits the metaphorical ascription of content may be real, and the ascription of content genuinely useful for some purposes, the content itself, if metaphorical, does not really exist; it is just a way of talking. So if you think that a 'red' quale is

a non-physical state, you have no reason to change your mind on discovering that it is correlated with metaphorical content. However, if perception is constituted by bodily interactions with objects in the environment, we can appeal to the qualities of the interaction to account for phenomenal qualities, and use metaphorical content, if appropriate, to explain how those interactions take place. This brings us to the other proviso, however, which is that content and representations may not be the most appropriate way to account for the physical interactions involved in perception, as we will see in the next section.

6.5 - Representation Restricted

According to Marr's (1982) theory, and accounts in the same tradition, conscious perception is realised by the brain's construction and deployment of representations. This processing is properly understood as an enabling rather than constitutive feature of perception. A constitutive account of perception says what perception is. An enabling account, in this context, could be thought of as a variety of causal account, where we take 'A causes B' to imply that A happens at the same time as B but at a lower explanatory level. Endorsers of the personal/subpersonal distinction observe that the representations in Marr's account must be enabling rather than constitutive features of perception because they are subpersonal, and perception is a personal-level faculty (McDowell, 1994b; Pessoa et al., 1998). Noë (2004) denies that the personal/subpersonal distinction is important, which means he cannot so straightforwardly establish that subpersonal representations are not constitutive features of perception. However, like Pessoa et al, he makes use of Marr's own distinction between the task and representational levels of analysis. The former specifies which tasks perception aims to perform, while the latter describes how the brain processes information, in the form of representations, to enable these tasks to be performed. As Noë appreciates, the distinction between the task and representational levels should be taken to track the constitutive/enabling distinction, meaning representations are mere enabling features even if the task level is not necessarily the personal level.

McDowell suggests that by appreciating the constitutive/enabling distinction, a rapprochement can be achieved between Marr's representationalism and ecological perception,

which is non-representationalist. His claim is that a Marr-style representationalist account can usefully be given of the subpersonal processes that enable perception even while we use the ecological approach to characterise perception's constitutive features at the personal level. As Pessoa et al. and Noë (2004) bring to our attention, however, Marr has a faulty view of representation that prevents this reconciliation from being achieved. The problem arises because Marr depends too heavily on his representational level framework to inform his constitutive or task level account of vision. Indeed, he appears to straightforwardly conflate the enabling with the constitutive features when he claims that: "vision is the *process* of discovering from [retinal] images what is present in the world, and where it is" (Marr, 1982, p. 2). The bluntest objection we could make to this is that it commits a straightforward category mistake. As McDowell has argued, vision is personal level process in which we come into contact with the world rather than the retinal image. Pessoa et al. are sympathetic to this point, but their argument does not focus on the bare coherence of claiming that perception is constituted by a subpersonal process. Their criticism is mainly directed at the danger that a subpersonal constitutive account invites implausible commitments including, as we saw in chapter 2, the undue commitment to a subperson, a homunculus, for whose benefit processes like filling-in need to take place.

Noë (2004), having denied that the personal/subpersonal distinction is important, has no call to insist that a constitutive account of vision must make reference to the whole person (or more broadly, agent). His relatively nuanced criticism, reflecting a point originally made by Nakayama (1994), is that Marr fails to appreciate the difficulty of giving an appropriate task level (i.e. constitutive) account of vision. As I take it, the problem is that Marr attempts to give an independently plausible account of perception's representational (i.e. enabling) features, and lets his view of the task level be unduly constrained by it, a bias implicit in his claim that perception *is* a process in which we use the retinal image to find out about the world. This is a mistaken order of priority, because our understanding of the processes that enable perception ought in many cases to be substantially constrained by independently motivated constitutive claims. Instead of only allowing one-way traffic from the enabling to the constitutive, or for that matter in the other direction, there should be a two-way interplay between our understanding of perception and our understanding of the processes that enable

it (see, McDowell, 1994b; Wheeler, 2013).

The objection is not that Marr's account, and Marr-style accounts, are framed as if no possible evidence, in particular empirical evidence, could disprove them. The problem is that their apparent plausibility as enabling accounts means that should they unduly inform or be conflated with constitutive accounts, they close off the possibility of more carefully examining perception's constitutive features. Re-emphasising neural representation's rightful place only as an enabling feature has a useful dual effect. It opens up conceptual space for a better constitutive account of perception. This, in turn, opens up the possibility that we must re-evaluate our claims about the processes that enable perceptual consciousness. This could have various effects, for instance constraining in previously unrecognised ways the kinds of content or representational vehicle that plausibly enable conscious perception, and, in principle, making it plausible that representation should be partially or even wholly dispensed with in favour of appeals to dynamic bodily interaction with the outside environment.

A proper constitutive account of perception ought to draw on a number of distinct methods and considerations. This includes empirical work. For example, the work on change blindness and inattention blindness referred to by O&N is enough to strongly indicate that the brain does not represent, all at once, a large and richly detailed visual field. A constitutive account should also pay careful attention to the phenomenology of visual experience, an aspect of sensorimotor theory most prominent in Noë's (2004; 2012) accounts. Noë claims that visual experience does not feel picture-like. This is, as I think, a straightforward and compelling aspect of Noë's in some ways enigmatic description of visual experience, exemplified by his claim that: "A perceptual experience doesn't analyze or break down into the experience of atomic elements, or simple features [...] the moment you stop and try to make a specific feature the sole object of your consideration – this shade of red, for example – it slips away from you in the sense that it exceeds what you can take in, in completeness, in an instant" (2004, p. 135).¹⁷ The fact that visual phenomenology is not pictorial constrains the shape that an enabling account ought to take. It offers one reason, for instance, to doubt that the character of perceptual awareness should be explained by pictorial representations in the

¹⁷ I offer a more detailed take on this point in chapter 7.

brain.

Noë suggests that attending to one's own visual phenomenology is enough to demonstrate that perceptual consciousness, instead of being pictorial, is constituted by the exercise of Sensorimotor Understanding. This claim has some appeal, but is hard to make go through purely on the basis of attention to one's own experience. A further argument is made by Noë (2010), who claims, as we saw earlier (section 4.3), that a perceiver would not be able to integrate what they consciously see with what they do if conscious vision did not involve Sensorimotor Understanding, or a similar egocentric understanding of their bodily alignment with objects in the outside environment. This can be understood as a logical, constitutive claim about vision, and places downward pressure on our understanding of the processing carried out in particular by the ventral stream to realise visual awareness. Notice that the point comes through even more sharply, however, when we think of it as an engineering question, namely the question of how the ventral stream achieves its function of helping the dorsal stream produce appropriate skillful bodily behaviour. The idea that the ventral stream must trade in representations of SMCs is a low-level claim about the processing that enables perceptual consciousness which we find, it turn, places upward pressure on a high level account of conscious perception's constitutive features. This is one case in which the sensorimotor theory makes productive use of an interplay between explanatory levels.

On the view of Sensorimotor Understanding I endorse, in which it is a feature of the perceiver's behavioural capacities, the claim that Sensorimotor Understanding is a constitutive feature of conscious perception could be supported by the need for ventral stream processing to integrate with dorsal stream processing too, especially if we endorsed the view that the ventral stream integrates with dorsal stream activity by representing possibilities for coarse-grained action (Ward, Roberts and Clark, 2011). Sensorimotor Understanding, on the behavioural capacity view, also makes a good fit with non-representationalist, dynamical systems accounts of perception, considering that they already appeal to sensorimotor capacities (Hurley, 1998; Noë, 2004; Chemero, 2009; Silberstein and Chemero, 2012; Burhmann, Di Paolo and Barandiaran, 2013).

A constitutive account of vision should also take account, I argue, of philosophically motivated worries about intentionality and phenomenal consciousness. Noë (2004) attempts

to do justice to McDowell's (1994a) claim that the content of perceptual experience is necessarily conceptual by suggesting that the bodily capacities we draw on in conscious perception are themselves conceptual. Noë (2012) retracts this claim and instead states that sensorimotor skills involve a form of nonconceptual content, while occurring on a continuum with conceptual skills rather than being radically distinct. I will not comment on the debate between conceptualists like McDowell and his opponents (e.g., Dreyfus, 2013; Evans, 1982), who claim that action and perception can be intentionally-directed without being fully conceptual. However, it is important to do justice to the basic insight that perceptual consciousness must be intentionally directed, and that this requires the perceiver to bring to bear some kind of skillful, active engagement with the world. It will not do to assume that content borne by neural representations can play this role, since it has not been convincingly established that neural representations bear genuine rather than metaphorical content. For this reason, appealing to neural states alone will not do at all. Accounts featuring content-bearing neural representations plausibly provide, in some cases, a useful framework for explaining the role the brain plays in producing behaviour, but it has not been established that they can, except by reference to behaviour, help account for genuine intentionality.

We might be tempted to deny that a naturalistic account of cognition needs to answer Hutto and Myin's metaphysical worries about the ascription of 'content', or that it is the job of a scientifically-oriented account of vision to explain how the capacities appealed to by non-naturalistic accounts of perception can be accounted for in physicalistic vocabulary. Consider, however, the purpose of the skill-based approach to phenomenal qualities, endorsed by Clark (2008) and on Clark's reading, with which I agree, by the sensorimotor account. The skill-based approach is meant to avoid the intuition that phenomenal consciousness is non-physical by identifying phenomenal consciousness with the exercise of skills. Its job is to avoid the dilemma I mentioned earlier (adapted from Ryle, 1979, by Silberstein and Chemero, 2012) in which phenomenal consciousness is either 'nothing but' brain activity or 'something else as well'. As I have argued, if the skills appealed to by the skill-based account are not sensorimotor skills, they depend on internal representation. If that representation is metaphorical, or a useful fiction (Sprevak, 2013), or anything other than robustly real, it cannot help answer metaphysically motivated, anti-physicalist intuitions about consciousness

better than appeals to nothing but brain state. This, in my view, is the biggest problem with attempts to defuse non-physicalist intuitions about qualia with appeal to representation.

One of the best reasons for taking Sensorimotor Understanding to be constituted by a capacity for skillful behaviour - rather than merely by knowledge of the sensory consequences of possible movement - is that it offers a promising way to account naturalistically for the intentional directedness, and by virtue of this the phenomenal character, of perceptual consciousness. Giving a proper account of the criteria a repertoire of behavioural capacities must meet to count as constituting mindedness, and sensorimotor understanding in particular, is not something I pursue here, and is something that would eventually need to be provided for this point to be established. However, it is hard to view the outward behaviour of a human perceiver and resist the intuition that they are intentionally-directed toward the features of the environment they interact with. This provides one motivation for regarding behavioural capacities as constitutive features of perceptual consciousness.

If we could be certain that neural representations cannot be genuinely contentful, then appeal to sensorimotor skills would appear to be the only hope for finding a naturalistic account of intentionality and so phenomenal character, and the case for regarding them as constitutive features of perception would be secure. As we have seen, Hutto and Myin offer strong reasons to doubt that subpersonal content is naturalisable by appeal to covariance or teleology. A possibility Hutto and Myin endorse, which I also endorsed in section 2.1.2, is that content could be adequately naturalised by appeal primarily to empirical evidence. In principle, it might be possible to routinely provide full physicalistic descriptions of the neural processes that underlie a conscious experience, and use a subpersonal theory that makes essential explanatory appeal to representation to provide bridging principles which make the relation between the neural processes and the conscious state intelligible. Without a conceptual solution to the 'hard problem' of content, it would remain hard to robustly justify the claim that neural states bear genuine rather than make-believe content, or in the language of McDowell (1994a), to say that the brain is responding in a person-like way to reasons and not just causes. However, the empirical success of the enterprise just described would provide strong grounds for abandoning the requirement for a naturalistic theory of perceptual consciousness to meet these strong metaphysical demands.

That such an account could be provided is an empirical bet. Although it is possible to provide representation or content-involving accounts of neural activity that provide an explanatorily useful framework for understanding brain activity itself, there is no guarantee that this subpersonal content can correspond reliably to the content or phenomenal character of perceptual consciousness (Noë and Thompson, 2004). The problem is not the epistemically-founded worry that it is difficult to intelligibly describe and explain the neural processes that yield perceptual consciousness unless we make reference to the contribution they make to sensorimotor coupling. The problem is that there may be no intelligible link between purely inner processes and perceptual consciousness. This point is evident in Hurley's (1998) defence of vehicle externalism, although we could accept the modest conclusion that an account of perceptual consciousness ought to make explanatory appeal to sensorimotor coupling without necessarily insisting on ECM, the claim that the substrate of perceptual consciousness actually includes the body and outside environment.

The question of whether perceptual consciousness depends on internal representation or a non-representational dynamical system depends on how well the respective approaches make sense of empirical work. However, there are aprioristically motivated reasons for not being surprised if it turns out to be impossible to explain perception by attributing content to brains or parts of brains, namely that intentionalistic vocabulary like 'believing' or 'representing' are words we apply in the first place to whole persons and not their brains. The fact that strategies like homuncular decomposition make it conceptually viable to apply these predicates to brains does not mean that we should expect these strategies to finally allow us to explain conscious perception by appeal only to contentful states of the brain. A better bet is to regard bodily capacities as perception's constitutive features and use these as the starting point to account for the physical processes that enable perceptual awareness.

7 - Object Experience as Temporal Experience

This thesis has so far focused largely on defending and clarifying the basic tenets of the sensorimotor theory. With these tenets shored up, there is as-yet untapped potential for the sensorimotor theory to explain in much greater detail the distinctive contribution it can make to accounting for specific perceptual phenomena. For instance, there is much work to be done on identifying the sensorimotor contingencies involved in colour perception, and the theory has not yet given much account of sense modalities other than vision. This present chapter continues the work clarifying and defending the sensorimotor theory, but also sketches out a novel contribution the sensorimotor theory can make to a long-discussed puzzle concerning the experience of temporally-extended events. In addition to explaining how the sensorimotor theory might make sense of event-experience, I will argue that object-experience may itself be a variety of event-experience, and that this would help make sense of Noë's appealing but oblique account of visual phenomenology.

7.1 - The Puzzle of Temporal Experience

The puzzle of temporal experience (Kelly, 2005) concerns perceptual awareness of temporally-extended phenomena like change, succession and constancy. To illustrate, Kelly offers the example of an opera singer delivering a sustained note. When you hear it, your experience incorporates not just the note's immediate presence, but also, somehow, its extended temporal duration:

There you are at the opera house. The soprano has just hit her high note – a glass shattering high C that fills the hall – and she holds it. She holds the note for such a long time that after a while a funny thing happens: you no longer seem only to hear it, the note as it is currently sounding [...] in addition, you also seem to hear something more [...] the note now sounds like it has been going on for a very long time [...] What you hear no longer seems to be limited to the pitch, timbre, loudness and other strictly audible qualities of the note. You seem in addition to experience, even to hear, something about its temporal extent. (Kelly, 2005, p. 208)

Vision, too, sometimes comprises the experience of temporally-extended events. Grush (2007)

offers an example of motion experience. Suppose you are looking at a clock. You could stare at the hour hand all day, and although you may infer that it is moving, you never visually experience its motion. When you look at the second hand, however, you seem to perceive its motion directly as it moves around the clock face. Temporal experience is a puzzle, as Kelly (2005) puts it, rather than a mere problem, because it involves an apparent paradox: What you perceive now, at this moment, ought only to be what is present now. Motion, change and duration are not momentary, but take place over extended periods. Solutions to the puzzle will differ in their diagnosis of what is happening when you hear the opera singer, or look at the second hand, but the examples clearly illustrate that there are data to be explained.

Dainton (2010) places proposed solutions into three distinct categories. The 'cinematic' approach is the view that perceptual experience is divided into static snapshots, like a cinema reel. On this approach, what you perceive at a given instant is just what is present at that instant, meaning you cannot right now literally experience a temporally-extended event. Temporal experience, here, has to be explained some other way. Crick and Koch (henceforth C&K) (2003), for example, suggest that you do not experience motion, but "a series of static snapshots, with motion 'painted' on them" (p. 122); they suggest, by analogy, that the experience of motion has something in common with the motion suggested by a drawing of a person in mid-stride (p. 122, figure 1). The other two types of explanation, by contrast, endorse the existence of what James (1890/1981) called the 'specious present', an experiential 'now' that really does have duration.

The 'retention-protention' approach (Dainton, 2010), of which Husserl (1991) and, under Husserl's influence, Varela (1999) were notable exponents, says that the experience of what is genuinely present at a given instant is accompanied by 'retentions' from the immediate past and/or 'protentions' into the future; these can be understood, respectively, as special perceptual sorts of memory and anticipation. Grush (2007) advocates a retention-protention account he calls the Trajectory Estimation Model. In support of the approach, he cites empirical work on a number of illusions, for example the 'cutaneous rabbit' illusion (Geldard and Sherrick, 1972, cited by Grush, 2007). The study involved subjects being given a sequence of taps on the arm, five at each position, in three successive positions along the arm. The subjects reported that instead of feeling the taps in just three different positions, they

experienced them as travelling in an even trajectory along the arm, ultimately extending beyond the location of the final tap. Because this trajectory would not have been experienced had the sequence been confined to the first five taps, Grush's suggestion is that an experience at a given instant can feature as part of its content prior occurrences (see also the colour-phi phenomenon, discussed, e.g., Dennett, 1991). Further, because the felt trajectory continued onward, Grush concludes that protentions – anticipations of future events – also figure in experiential content.

The 'extensional' approach (or 'extensionalism') – endorsed, for instance, by Dainton (2000) – says that the content of the perception (the state of affairs represented by the experience) temporally tracks the vehicle (the physical state that realises the experience). Dainton (2010) observes that on this view, since “our episodes of experiencing are themselves temporally extended, [they] are thus able to incorporate change and persistence in a quite straightforward way”.

Clark (2006) uses temporal experience – in particular, Kelly's opera singer example – to expose what he regards as a general problem with the sensorimotor theory. His argument asserts that temporal experience poses a problem, in particular, for Noë's notion of presence-as-access, the view that you fail to experience the objects of perception as straightforwardly 'present', but instead experience them as 'accessible', thanks to your possession of Sensorimotor Understanding. In targeting presence-as-access, Clark implicitly targets the claims that perceptual experience depends on Sensorimotor Understanding and the claim that phenomenal character can be characterised by reference to SMCs. Clark says we cannot explain the experience of the opera singer's note “by appeal to any sense of the potential availability of the missing parts of the temporally-extended sound stream, nor can we know (indeed, it is barely intelligible to ask) how those missing parts of the soundstream would vary or come into focus as we move our head or body” (Clark, 2006, p. 23). Since the past and future are not, in other words, mediated now by laws of sensorimotor contingency, it cannot be Sensorimotor Understanding that explains the experience of the note's duration.

Noë (2006) responds by agreeing that Sensorimotor Understanding only explains object perception, arguing that this does not compromise the sensorimotor theory in general. He suggests that since event perception is, in any case, a quite different species to object

perception, the theory is not committed to giving a sensorimotor account of the opera singer case. He justifies this by reference to the distinct qualitative character of the temporal experience compared to the experience of an object: “it rides roughshod over the phenomenology [...] to say that the past sounds [like objects] are now present or that they are now accessible” (p. 28); instead, he claims, you hear the note “as having a certain trajectory or arc, as unfolding in accordance with a definite law or pattern” (p. 29).

Noë’s positive account of the opera singer case posits that we have a (non-sensorimotor) grasp of where the note is coming from, analogous to linguistic understanding:

When you hear the singer’s sustained note, you do not experience the acoustical properties of the sound, any more than you experience the acoustical properties of the words you hear when you understand speech. In the linguistic case, you hear meanings themselves, you hear what is said. In the case of the singer, what you actually hear is the singer herself, her voice, her vocal action – what she is doing. It is the fact that the singer is doing something, performing an action, that fixes the relevant temporal horizon and intentional arc (Noë, 2006, p. 29)

This offers a solution to the temporal experience puzzle akin, in important respect, to C&K’s cinematic account. Where Noë claims that different kinds of knowledge are responsible for object experience and temporal experience, C&K, similarly, ascribe these to separate “mechanisms” (p. 122). Moreover, both C&K and Noë agree that you do not experience temporally-extended phenomena directly: As with Noë’s invocation of an experienced trajectory, C&K suggest motion is “painted” on” (p. 122) to static snapshots, resulting in your experiencing moving objects as being in motion, while failing to directly experience that motion.

However, Noë maintains that even temporal experience involves a coupling with the environment rather than a representing:

[P]erception is an activity of sensorimotor coupling with the environment [...] experiences are not acts [...]; they are not representations; they are activities, events themselves; they are temporally extended patterns of skillful engagement. When you perceive an event unfolding, it is not as if you occupy a dimensionless point of observation. You live through an event by coupling with it. (Noë, 2006, p. 31)

This suggests a quite different line of response than the one endorsed by the analogy with linguistic understanding. Extrapolating from Noë, the thought seems to be that the physical substrate of the experience is a smoothly continuous activity rather than one which breaks down into temporally discrete chunks: hence to explain the physical substrate of perception, one must look at dynamically unfolding interactions, rather than ‘object’-like structures in the brain. If we appended to this the extensionalist view that the content of an experience temporally tracks the activity of experiencing, it would mean that the content of experience is always, itself, temporally-extended. Notably, this does not need to entail that past and future portions of the note should be experienced as now present. Rather, it suggests that to experience the world is – as Noë claims – to experience being in the midst of some trajectory.

7.2 - Coupling and Knowing in Temporal Experience

The argument over temporal experience speaks to a broader tension within the sensorimotor account. As many commentators have noted, the approach sometimes appears to stipulate that the temporally-extended bodily exercise of sensorimotor skill is required; at other times, it apparently suffices that the perceiver possesses sensorimotor skill. Aizawa (2010) describes these as ‘strong’ and ‘weak’ versions of the theory. Noë’s discussion of temporal experience betrays a similar tension: while its main claim is that the experience of duration is explained by your non-bodily understanding of what you are hearing, it indicates a quite different line of response when it argues that hearing the note involves, in any case, a temporally-extended coupling with the environment. The tension is particularly stark in the case of event perception, because it is hard to see why temporal experience should be explained both by understanding and sensorimotor coupling. In the original story about object perception, sensorimotor knowledge is sometimes glossed as knowing how to act, for example when O&N suggest sensorimotor knowledge is comprised of “action recipes” (p. 945) or when Noë says: “[t]o experience [an object] as on the left is to experience it as necessitating [...] various possibilities of sense-affecting movements” (2004, pp. 87–88). Here, we can readily grasp that perception might involve, vitally, a temporally-extended process of bodily coupling, itself

featuring Sensorimotor Understanding. By contrast, your knowledge of where the opera singer's note is coming from may be implicit, but there is no obvious way in which it is practical, or geared toward action, as sensorimotor mastery might be. If it is not practical, but, as Noë suggests, more like linguistic comprehension, then it is hard to see why possessing or exercising this knowledge should entail, in any interesting sense, a coupling with the environment rather than just a representing of the environment. The sensorimotor theorist, in the temporal case, no longer has an obvious response to the theorist who maintains that neural states or structures alone are the interesting, indeed constitutive, features underlying conscious perceptual experience.

This threatens to undermine the sensorimotor theory, as it means the 'weak' variant, on the current account of temporal experience, can no longer sustain any of the theory's main tenets. Vehicle externalism or ECM is ruled out, since bodily movement is not required for perceptual experience. Sensorimotor Understanding, SMCs and presence-as-access only give an incomplete account of perceptual experience because they account only for object experience, not temporal experience. The sensorimotor theorist could respond that although event perception is non-sensorimotor, it depends for its existence on object perception, which is sensorimotor. Conceding this much, however, gives the opponent room to deny that perception is intrinsically sensorimotor at all: Clark (2008), as we have seen, claims that perception is a matter of sensorimotor summarising – the extraction of information about sensorimotor contingencies, along with other information, for the construction of representations that are not themselves finely sensitive to the sensory effects of possible movements.

Moreover, the distinction made by Noë between events and objects seems tenuous. Noë (2006) says "objects [unlike events] are timeless in that they exist whole and complete at a moment in time" (p. 28). The implication is that it is therefore reasonable to suppose that *perception* of objects and events are different matters. On this view, perception of a moving object would seemingly be a matter of perceiving the object (sensorimotor) and perceiving its trajectory, an event (non-sensorimotor). However, if you perceive an unmoving, unchanging object, and perceive it to be constant, you expect that it existed moments before, and that it will continue to exist moments into the future. As a result, it is not enough to tack on some

kind of non-sensorimotor perception just for some instances. The non-sensorimotor element would apply to all instances of perception.

The opera singer example is a little misleading in this respect, because the phenomenal experience of the note's temporal duration seems like an optional extra, a side issue to the perception of immediately audible qualities like pitch, loudness, and so forth. Temporal aspects to perception are not usually like this. When you cross a busy road, or catch a ball, the perception of objects as still or in motion is no side issue, but central to your engagement with them. For this reason I think it is better to hold that object and event perception are not different types. The fuzziness of the distinction between object and event perception provides another reason to prefer a unified picture for objects *and* events, and conceding that events are not perceived in sensorimotor fashion provides a reason to believe that objects are not perceived in such a fashion either. This is one reason why I will argue presently that a unified sensorimotor model can be offered for both.

The best solution, I suggest, is to drop Noë's analogy with linguistic understanding and stick, instead, to an extensionalist story, which says that event experience is explained by a particular kind of temporally-extended coupling with the environment. This requires adopting the 'strong' sensorimotor theory, in the sense defined above, and accords with my suggestion in chapter 4 that Sensorimotor Understanding is a criterion that temporally-extended engagements must meet, rather than an entity that enables the appropriate engagements to take place. The mere possession of sensorimotor skill would not be sufficient to explain the experience of duration, since duration is not a matter of sensorimotor contingency. However, the bodily exercise of Sensorimotor Understanding takes time anyway: so, if we assume the content of experience temporally tracks the vehicle, the experience of duration comes for free. This suggestion is not merely a get out clause, but fits the phenomenology – the experience of trajectory – aptly described by Noë. It explains why your experience, now, of the opera singer's note sounds like it is part of something temporally-extended.

7.3 - Object Experience and Skill-Based Access

I have suggested that the best response the sensorimotor theorist can make to Clark's objection from temporal experience is to argue that the experience of temporal duration supervenes on a temporally-extended physical process, comprised of skill-driven bodily coupling with the environment. This means that the exercise of Sensorimotor Understanding must be temporally-extended in the case of event experience. The sensorimotor theory could, conceivably, endorse this, but claim that object experience does not depend on temporally-extended engagements, merely on the activation of representations of SMCs. In this case, the debates over temporal experience and what Aizawa calls 'weak' and 'strong' variants of the sensorimotor theory come apart. My proposed version of extensionalism, however, accounts for object perception as well as event perception. This results from what I contend is the most productive way to understand presence-as-access, and it is this thesis I turn to now.

As we saw earlier, Noë (2004) outlines his notion of skill-based access when he likens perception to accessing a newspaper via the World Wide Web. He points out that when you view the online version of the New York Times, your computer does not download the day's edition all at once, but downloads, on request, one article at a time. This is sensible, as it limits the burden placed on your computer and internet connection, and means that should an article be updated, you get the latest version. Importantly, accessing the paper in this piecemeal fashion is, for all intents and purposes, just like having the whole issue there at once, since every article is accessible as needed. The day's edition is available, as Noë puts it, *virtually*.

Seeing, he suggests, works in a similar way. You do not experience, all at once, a richly detailed visual field. Instead, you access, as required, detail available from the outside environment. The presence of rich detail is also, in this sense, virtual (Noë, 2004, pp. 49–51). However, in a crucial disanalogy with the computer case, Noë claims that perceptual presence is "virtual *all the way in*" (2004, p. 134). The suggestion is apparently that no complete property is ever present to experience in a local or offline manner: a claim Noë argues is justified by attention to one's own experience:

A perceptual experience doesn't analyze or break down into the experience of atomic elements, or

simple features [...] the moment you stop and try to make a specific feature the sole object of your consideration – this shade of red, for example – it slips away from you in the sense that it exceeds what you can take in, in completeness, in an instant. (2004, p. 135)

This is a puzzling sort of claim. There is (I think – and the reader might accept, at least for the sake of argument) something in Noë’s phenomenological claim that you cannot, in a momentary act, get a visual grasp on any complete property, such as a shade of red. However, there is certainly something it is like to see a red thing. Noë explains this by claiming that properties are “present not as represented, but as accessible [...] [t]hanks to my possession of sensorimotor skills” (2004, p. 215). As he puts it in a later piece, the environment:

[...] shows up as present, but out of view, in so far as I understand that I am now related to it by familiar patterns of motor-sensory dependence. It is my basic understanding of the way my movements produce sensory change given my situation that makes it the case, now, even before I have moved an inch, [emphasis added] that elements outside focus and attention can be perceptually present. (Noë, 2012, p. 19)

This view is a ‘weak’ position, since it stipulates that no movement or temporal extension is required. As a result, the perceptual experience described must be the result of an internal state, perhaps an internal representation specifying what movements the perceiver should make to bring objects into view. It is not obvious, however, why this sort of Sensorimotor Understanding should not allow a perceiver at some instant to take in an atomic visual feature, given visual features are meant to be specifiable by reference to movement-related contingencies. We could make sense of this by supposing that the sensorimotor mastery made use of by the perceiver at a given moment fails to correspond, precisely, to any atomic feature, such as a shade of red. If this were true, however, it would not be clear how we can explain the fact that perceivers experience shades of red at all. As I result, I contend that the best way to make sense of the computer metaphor, and the associated phenomenology, is to hold that object experience, like event experience, depends on a temporally-extended physical process. Indeed, the extensionalist proposal is useful as a way to characterise both event and object perception. I elaborate on this in the next section.

7.4 - Extensionalism About Event and Object Experience

It is beyond my present scope to consider the advantages and disadvantages of various approaches to temporal experience independent of their relation to sensorimotor theory. I argue, however, that extensionalism is likely to provide the right conceptual foundation for a sensorimotor account of temporal experience. Adopting this approach is useful for the sensorimotor theory, because it allows the theory to respond convincingly to the specific challenge set by Clark (2006), and thereby save the approach from the danger that it fails altogether because it is unable to account for temporal experience. More significantly, there is good reason to endorse the extensionalist claim that perceptual experience temporally tracks a temporally-extended physical process anyway, since it has the added utility of offering a new and better way of explicating some of the compelling claims featured in Noë's account of presence-as-access. On my proposal, event and object experience are not only explained the same way, but are aspects of the same phenomenon.

An extensionalist sensorimotor account, in the first place, explains temporal experience by claiming that perceptual awareness of a temporally-extended event supervenes on a temporally-extended process of interaction between the perceiver and the environment, in such a way that the content temporally tracks the vehicle. This is, in principle, compatible with the view that object experience is explained by the mere activation of a representation bearing content about SMCs. However, we could plausibly go further and suppose that the characterisation of object experience found in Noë's claims about presence-as-access can best be explained by accounting for object experience in the same way as event experience.

An initial move is to reject any suggestion that knowing the movements you could make to bring an object into view can amount, by itself, to entertaining perceptual content about that object. Instead, knowing this might enable, more modestly, a nonspecific feeling that there is detail available which can be accessed from the environment as needed (as suggested by Clark, 2008, p. 194). To account for perceptual presence – of an opera singer's note, or a shade of red – I suggest we instead point to the perceiver's skillful coupling with an environment in which the relevant laws of sensorimotor contingency apply. Suppose, for the sake of argument, that this is a correct account of object experience. Combined with the

extensionalist claim that the content of perceptual experience temporally tracks the physical substrate, it would follow that the content of object experience is itself temporally-extended. This is counterintuitive, but compelling if you consider the presence-as-access in a particular light. Recall two of the claims featured in this aspect of the sensorimotor account.

There is a phenomenological claim, concerning a perceiver's inability to experience, in a momentary act, any atomic visual features. There is also a claim about the material substrate of experience, captured by Noë's discussion of virtuality, which suggests that perception is entirely beholden to the agent's online interaction with the environment. Neither of these claims is indisputable, but they are coherent accounts of how perception may work. If you think there is something right about the theses, extensionalism about object experience offers a way to make sense of it. On the view I recommend, one cannot, in a momentary act, take in a shade of red for the very same reason that one does not experience, at this instant, the past and future portions of the opera singer's high C. Perceptual experience, both of objects and events, supervenes on something that may, essentially, be understandable only by reference to a temporally-extended, dynamical process.

Noë's (2004) claim that experience is "virtual *all the way in*" (p. 134) could suggest that the spatial size of the visual field is infinitesimally small or even non-existent. This makes the sensorimotor account sound like it denies perceptual experience altogether. A better gloss on the claim says that experience is virtual all the way in not, primarily, because the information processed at one time by the visual system is minimal – although it may be minimal – but because this information can only be processed in the course of temporally-extended coupling with the outside environment. To revisit the World Wide Web analogy: if your access to the online newspaper were virtual all the way in, a faulty understanding of virtuality would suggest that your computer 'knows' how to access the internet, but has no screen with which it can locally display a web page. According to my gloss, however, the computer has a screen, but loses its ability to display any web page – even one you already have open – the instant it loses its connection to the Internet. It is a commonplace that perception, ordinarily, requires that there is an environment present to perceive. However, the point of virtuality, as currently understood, is to show that ongoing dynamic interaction with the environment is a conceptual necessity for perceptual experience.

An upshot of this view is that perceptual experience can only supervene on a temporally-extended process, and never on a momentary state. This goes part way to accounting for the experiential quality that Noë observes accompanies presence-as-access, namely the inability to fixate, at an instant, on any atomic visual feature. Note, however, that the content of an experience need not temporally track the material realiser. If it does not, a temporally-extended process could yield the experience of an instant in time in which you take in a shade of red. The extensionalist view that the content does temporally track the vehicle has the benefit of offering an apt way to reformulate the phenomenological account offered by Noë. The inability to take in a shade of red in what you experience as a durationless instant is explained by the fact that having an experience of red is, in part, having an experience of duration. A visual feature never feels phenomenally present at this instant because, before you have a chance to fixate on it, the relevant detail has lapsed from being something that you are confronted with now, to something that you have been confronted with a moment ago. Extensionalism, here, allows that you can experience a shade of red, but stipulates that this takes place – and is experienced as taking place – over an interval of time. The process is likely to be assisted if your movements during that time help you gain information about the patterns of sensorimotor contingency that currently hold between your body and the environment.

Returning now to the debate between Clark and Noë: Clark's objection to presence-as-access is that elapsed portions of the opera singer's note are not accessible now, hence your experience of them cannot be explained by means of their accessibility. Noë's response denies that object and event experience need to be explained in the same way, on the grounds that they are different species of awareness: "objects are primary in our experience [...] experience of events depends on a more basic sensitivity to the presence of objects" (Noë, 2006, p. 31). I earlier argued that Noë's line of response may be inadequate to defend the central tenets of the sensorimotor theory. This drawback can be avoided, I suggest, precisely by reversing Noë's claim, and contending that the experience of objects depends on the ongoing conscious presence of events that have been occurring moments into the past.

7.5 - Chapter Conclusion

It is a matter of some importance that the sensorimotor approach has the resources to provide a plausible answer to the puzzle of temporal experience. Noë (2006) offers an ambivalent response. He suggests, briefly, that it can be explained by temporally-extended ‘coupling’. As I have argued, however, there are no obvious grounds to endorse this, given his main claim that the experience of the opera singer’s note is explained by an understanding, akin to linguistic comprehension, of where the note is coming from. While this latter claim provides an account of temporal experience, it entails an abandonment of a number of the sensorimotor theory’s central tenets, as far as event experience goes. In so doing, it renders the theory more vulnerable to the objection that object perception is not fundamentally sensorimotor either, expressed, for example, by Clark (2008), when he espouses sensorimotor summarising.

With the aim of defending sensorimotor theory, I have aimed to indicate the conceptual foundation of a sensorimotor enactivist account of temporal experience. This involves rejecting the comparison with linguistic understanding and emphasising temporally-extended activity. Adopting extensionalism as a working hypothesis invites, in turn, a particular understanding of object experience. This is worth endorsing in its own right because of the useful gloss it places on Noë’s otherwise compelling ‘virtuality’ metaphor, and the phenomenology of presence-as-access. This consideration gives the sensorimotor theory even more reason to pursue extensionalism as an explanation of temporal experience.

Clark’s (2006) critique has the useful effect of forcing the endorser of sensorimotor theory to evaluate carefully the respective roles of coupling and knowing in the approach. If my suggestion is right, understanding only plays a role if it is directly implicated in bodily coupling with the environment. If bodily interaction is key, then temporal extension is also surely vital. The sensorimotor theory should, therefore, do more than pay lip service to temporality, and take into serious account the temporally-extended nature both of perceptual experience and its material substrate.

8 - Thesis Conclusion

In spite of all the attention the sensorimotor theory has been given since its first official statement, many of its tenets have continued to be incompletely or ambiguously defined, and one of my goals in this thesis has been to go some way toward clarifying them. This undertaking is closely related to the separate project of responding to some of the criticisms that have been levelled against the theory. I have focused in particular on Clark's argument that perceptual experience is likely to depend on the activation of representations which are not intimately involved in bodily movement and which stand-in for types, categories and relative locations rather than laws of sensorimotor contingency. To the end both of responding to Clark's argument and arguing for a particular understanding of the sensorimotor theory, I have suggested that there are good reasons to think that perceptual experience depends on the exercise of behavioural capacities, which may be relatively coarse-grained. In this final passage, I am going to tie together some key strands from the main part of the thesis.

Sensorimotor contingency is perhaps the most important concept deployed by the sensorimotor theory. However, there is sometimes equivocation about what SMCs actually are. One potential tension, although I do not find it particularly problematic, is the shift of emphasis in Noë's sole-author accounts from SMCs governing subdoxastic sense inputs to the SMCs that respectively define P-properties and objective properties, the latter being characterised by the ways P-properties change with movement. More could be said about the relation between these different kinds of property. However, there is nothing obvious to suggest that a final version of the sensorimotor theory could not refer to a combination of all these kinds of sensorimotor contingency.

More problematic is the slide many statements of the sensorimotor account make between describing SMCs as counterfactual conditionals, i.e. the results of merely possible movements, and as patterns of bidirectional output-input dependence manifest in the agent's sensorimotor engagements right now. This distinction is similar to the one suggested by Burhmann, Di Paolo and Barandiaran's concepts of 'sensorimotor environment' and 'sensorimotor habitat'. For the sake of conceptual hygiene, I think SMCs should be

understood only to describe the results of counterfactually possible movements, meaning it need not be possible to identify them from the agent's present bodily interactions. SMCs, on this understanding, describe the properties a subject exhibits an appropriate kind of sensitivity to if she perceives, but do not presuppose that the subject does perceive. The concept, so construed, avoids prejudging the means by which subjects exhibit the appropriate sensitivity, which could involve the activation of representations, the exercise of behavioural capacities, or both.

Further, adopting the counterfactual conditional understanding of SMCs does justice to the sensorimotor theory's central suggestion that perceivers, by virtue of possessing sensorimotor mastery, experience the presence (in absence) of parts of the environment even when they are not presently interacting with those features in an outward bodily way. I have sometimes heard people suggest that SMCs are properties a perceiver can 'have'. This is inaccurate. They could be properties a perceiver's interactions with the environment have. On the understanding I prefer, they are bodily relations a perceiver stands in to the outside environment which we describe by specifying what would happen *if* the subject or objects in question moved in particular ways. One reason it is important to clearly understand the concept of sensorimotor contingency is that a next phase for the sensorimotor theory will be to give detailed descriptions of the SMCs that comprise specific experiential properties, such as shades of red. Moreover, how one understands the concept of sensorimotor contingency will have a knock on effect on how one understands the concept of sensorimotor mastery as well as the physical mechanism implicated in it.

As I suggested earlier in the thesis, sensorimotor mastery admits two possible directions of fit. It is sometimes glossed as Being Able To Act, which implies a desire-like direction of fit, and sometimes as Knowing What Would Happen If You Did Act, which implies a belief-like direction of fit. One proposal is that sensorimotor mastery is essentially belief-like, and that having an ability to act is merely a trivial consequence of possessing sensorimotor mastery. Noë offers this interpretation at times, although not consistently. The construal of sensorimotor mastery in which it is primarily belief-like invites the suggestion that the exercise of sensorimotor mastery is constituted by the deployment of internal representations bearing content about SMCs. The representations could take the form, for

instance, of the predictive models proposed by Seth (2014). Although this sort of account could turn out to be correct, I do not think endorsing it is the sensorimotor theorist's best bet, since the proposal entails abandoning many of the improvements the sensorimotor theory otherwise makes on the orthodox approach.

One improvement the sensorimotor theory would otherwise make is to avoid placing an unsustainable explanatory burden on the concept of internal representation. There are significant limitations to the explanatory efficacy of appeals to internal representation. Hutto and Myin (2012) have demonstrated that it is not clear that subpersonal content can ever be accounted for in a naturalistically respectable way. Even if one were not moved by Hutto and Myin's contention that any and all content-talk should therefore be avoided, there are sensible limits that should be placed on internal representation's theoretical role. Since it has not been established that neural states can bear metaphysically real content rather than content we ascribe as a useful metaphor because it is explanatorily expedient to do so, we should not appeal to content to account for phenomenal qualities, particularly if we, as physicalists, hope to respond to those with metaphysically-motivated anti-physicalist intuitions about qualia. Setting this point to one side, it is not certain, in any case, that appealing solely to the content and functional role of neural representations will ever be sufficient by itself to make it intelligible how personal level perceptual content or perceptual phenomenology is connected to the physical.

These drawbacks also speak against the objection, from Clark (2008), that perceptual experience depends on the activation in the ventral stream of internal representations geared towards the exercise of epistemic skills like sorting, sifting and classifying, and that the sensorimotor theory is therefore implausible. If these skills are not bodily skills, the view ends up depending on the content of the internal representations to account for phenomenal qualities and, as I have suggested, internal representations do not appear to be up to this job. If sorting and sifting are varieties of sensorimotor skill, as is plausible, then Clark's account of ventral stream perception does not challenge the sensorimotor theory's core claim that perception is constituted by the exercise of sensorimotor skills, which we can allow are likely to be coarse-grained rather than fine-grained.

Sensorimotor theorists also have a strategic reason not to endorse the claim that the

exercise of sensorimotor understanding is constituted by the activation of representations, since doing so makes the theory a hostage to fortune. If the phenomenology of presence-in-absence and the logical relation between perception and action could be compellingly accounted for in a way that no longer entails that perceivers master SMCs, then few of the sensorimotor theory's core tenets would remain secure. Having said that, if it were established the perceptual experience does depend on nothing more than the activation of internal representations, the sensorimotor theory would continue to have promise as an account of phenomenal qualities. Internal representations, supposing that they do bear genuine content, could not bear the hyperintensional content required to actively discount the possibility that representations of types, categories and locations are also representations of SMCs. Appealing to SMCs rather than types, categories and locations could, however, help representationalists about phenomenal qualities respond to the Inverted Earth counterexample.

Although there are problems with tying the sensorimotor theory too heavily to the concept of internal representation, sensorimotor mastery must, at least in a loose sense, involve Knowing What Would Happen If You Did Act, since SMCs describe the consequences of merely possible movements, and perceivers must grasp them in some way. My suggestion is that sensorimotor mastery nonetheless does not depend essentially on internal representation, but on behavioural capacities. These capacities, I suggest, implicate Knowing What Would Happen If You Did Act just in case they meet appropriate criteria. This approach perhaps comes close to what Noë would say is the behaviourist mistake of conflating effect (i.e. behaviour) with cause (i.e. knowledge). However, it avoids crude behaviourism. First, the idea is not simply to abandon sensorimotor mastery as a belief-like capacity and appeal to behavioural capacities instead. My suggestion is that the relevant behavioural capacities logically implicate the appropriate understanding, meaning that the understanding is a real entity even if it is not reified as something like a neural representation. Thus it does not deny perceptual consciousness, but explains perceptual consciousness by appeal to sensorimotor mastery. Further, the present view avoids the implausible suggestion that perceivers must in every case move their bodies in order to perceive, since the behavioural capacities may be active merely by virtue of a subject's being *prepared* to behave. Preparedness

to behave could, in principle, be ascribed on the basis of neural and bodily states without appeal to representations or presently-occurring bodily movements.

Finally, the view I am recommending does not rule out the possibility of giving a restricted theoretical role to internal representation. In fact, it endorses pluralism about how the relevant behavioural capacities may be accounted for physicalistically. It is compatible with an entirely nonrepresentational dynamical systems approach. It is also compatible with an approach that appeals to information processing, including the deployment of internal representations, should they prove the most effective way to explain how the appropriate behavioural capacities are realised. This approach to representation avoids unduly conflating the relatively innocent suggestion that representation is a useful metaphorical posit for making sense of the physical processes implicated in perception with the faulty suggestion that perception just is the activation of an internal representation.

A final thing is to see how the behavioural capacity view can make sense of the phenomenology of object experience described by Noë, as well as provide a distinctive answer to the puzzle of temporal experience, which present a specific challenge to the sensorimotor theory as well as presenting a challenge to accounts of perception more generally. This is my suggestion: Perception unfolds as a temporally-extended process involving the exercise of capacities for bodily behaviour. This process frequently involves outwardly observable bodily movement, although it need not, since behavioural capacities are often exercised merely in the course of becoming prepared to move our bodies. This process cannot be broken down, in a principled way, into discrete time slices. Considered at any particular instant, however, a perceiver is poised to make certain bodily movements. This poise reflects the perceiver's understanding of certain SMCs, where SMCs are taken to be characterised by counterfactual conditionals specifying the results of possible movements. The movements a perceiver is poised to make at a particular instant are sufficient to account for the perceiver's experiencing the presence in absence of certain environmental detail. However the SMCs they are exercising mastery of at this instant are insufficient to determinately specify any particular environmental property, for example a shade of red.

The experience of a determinate feature, such as a shade of red, only happens in the course of the temporally-extended exercise of that behavioural capacity. Over a period of time,

a perceiver undergoes a continually evolving poise to make skillful bodily movements. Over the course of this evolution, the perceiver may experience a shade of red, although they can never fixate on the shade of red, since seeing the shade of red depends on this evolving, temporally extended, non-atomic process. If it the content of the experience did not temporally track the vehicle, there need not be any obstacle to seeing the shade of red all at once, as a non-temporally extended content could supervene on a temporally-extended vehicle. But we can make sense of Noë's description of visual phenomenology, which I find compelling, by supposing that the content temporally tracks the vehicle. This means the experience of the shade of red is an experience of temporal extension. This is suggestive of a different response to the opera singer case than the one given by Noë, namely a version of extensionalism rather than the cinematic approach he endorses. This gives us a way to account for the phenomenology associated both with event and object experience.

One thing I have not addressed in detail in this thesis is the precise mechanism implicated in sensorimotor mastery, and this aspect of the sensorimotor theory continues to be in need of significant development. Further, I have not examined in detail the precise character of the behavioural capacities that would implicate different forms of sensorimotor mastery; offering such an account has the potential to be a useful next step for the sensorimotor theory. A final noteworthy limitation is that I have not resolved the more general question of how to licence the claim that appropriate behaviours and behavioural capacities are themselves intentionally-directed. Answering this question might involve endorsing the biological enactivists' appeal to autonomy and normativity, although we might finally decide that notions like these are dispensable. This is a question that the sensorimotor theory ought to take a stance on. However, I believe that with a clear understanding of the fundamentals of the sensorimotor theory, and an appreciation of its significant merits, we ought to be well placed to progress the theory further.

Bibliography

- Aizawa, K. (2010). Consciousness: Don't give up on the brain. *Royal Institute of Philosophy Supplements*, 67, 263–284.
- Allen, C. (2009). Teleological Notions in Biology. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2009).
- Ballard, D. H., Hayhoe, M. M., Pook, P. K., & Rao, R. P. N. (1997). Deictic codes for the embodiment of cognition. *Behavioral and Brain Sciences*, 20(4), 723–742.
- Bennett, M. R., & Hacker, P. M. S. (2003). *Philosophical Foundations of Neuroscience*. Malden, Mass.: Wiley-Blackwell.
- Berlin, B., & Kay, P. (1969). *Basic Color Terms: their Universality and Evolution*. Berkeley and Los Angeles: University of California Press.
- Bermudez, J. L. (2000). Personal and subpersonal: A difference without a distinction. *Philosophical Explorations*, 3(1), 63–82.
- Blackmore, S. J. (2002). There is no stream of consciousness. *Is the Visual World a Grand Illusion? Spec. issue of Journal of Consciousness Studies*, 9(5-6), 17–28.
- Block, N. (1990a). Consciousness and accessibility. *Behavioral and Brain Sciences*, 13(4), 596–598.
- Block, N. (1990b). Inverted Earth. *Philosophical Perspectives*, 4, 53–79.
- Block, N. (1996). Mental paint and mental latex. *Philosophical Issues*, 7, 19.
- Block, N. (1998). Is experiencing just representing? *Philosophy and Phenomenological Research*, 58(3), 663–670.
- Block, N. (2001). Behaviorism revisited. *Behavioral and Brain Sciences*, 24(05), 977–978.
- Block, N. (2005). Review of Alva Noë, *Action in Perception*. *Journal of Philosophy*, 102(5), 259–272.
- Brooks, R. (1991). Intelligence without representation. *Artificial Intelligence*, 47, 139–159.

- Buhrmann, T., Di Paolo, E. A., & Barandiaran, X. (2013). A dynamical systems account of sensorimotor contingencies. *Frontiers in Psychology*, 4.
- Burge, T. (1979). Individualism and the mental. *Midwest Studies in Philosophy*, 4(1), 73–122.
- Campbell, J. (2002). *Reference and Consciousness*. Oxford: Oxford University Press.
- Campbell, J. (2008). Sensorimotor knowledge and naïve realism. *Philosophy and Phenomenological Research*, 76(3), 666–673.
- Carruthers, P. (2011). Higher-order theories of consciousness. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2011.).
- Chalmers, D. J. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. Oxford: Oxford University Press.
- Chemero, A. (2009). *Radical Embodied Cognitive Science*. Cambridge, Mass.: The MIT Press.
- Chemero, A., & Cordeiro, W. (2000). Dynamical, Ecological Sub-Persons. *A Field Guide to the Philosophy of Mind*.
- Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Ryu, S. I., & Shenoy, K. V. (2010). Cortical preparatory activity: Representation of movement or first cog in a dynamical machine? *Neuron*, 68(3), 387–400.
- Clark, A. (1997). *Being There: Putting Brain, Body, and World Together Again*. Cambridge, Mass.: The MIT Press.
- Clark, A. (2006). That lonesome whistle: A puzzle for the sensorimotor model of perceptual experience. *Analysis*, 66(289), 22–25.
- Clark, A. (2008). *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford: Oxford University Press.
- Clark, A. (2009). Spreading the joy? Why the machinery of consciousness is (probably) still in the head. *Mind*, 118(472), 963–993.
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(3), 181–204.
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58(1), 7–19.

- Clark, A., & Toribio, J. (2001). Sensorimotor chauvinism? *Behavioral and Brain Sciences*, 24(05), 979–980.
- Crick, F., & Koch, C. (2003). A framework for consciousness. *Nature Neuroscience*, 6(2), 119–126.
- Dainton, B. (2000). *Stream of Consciousness*. London: Routledge.
- Dainton, B. (2010). Temporal Consciousness. *The Stanford Encyclopedia of Philosophy, Fall 2010 Edition*. Internet Encyclopedia.
- Davidson, D. (1987). Knowing one's own mind. *Proceedings and Addresses of the American Philosophical Association*, 60(3), 441–458.
- Degenaar, J., & Myin, E. (2013). The structure of color experience and the existence of surface colors. *Philosophical Psychology*, (3), 1–17.
- Dennett, D. C. (1978). *Brainstorms: Philosophical Essays on Mind and Psychology*. Cambridge, Mass.: The MIT Press.
- Dennett, D. C. (1987). *The Intentional Stance*. Cambridge, Mass.: The MIT Press.
- Dennett, D. C. (1991). *Consciousness Explained*. Boston: Little, Brown and Co.
- Dennett, D. C. (1992). "Filling in" versus finding out: A ubiquitous confusion in cognitive science. In P. Van de Broek & D. C. Knill (Eds.), *Cognition: Conceptual and Methodological Issues* (pp. 33–49). Washington, DC: American Psychological Association.
- Dennett, D. C. (2001). Surprise, surprise. *Behavioral and Brain Sciences*, 24(05), 982–982.
- Dennett, D. C. (2002). How could I be wrong? How wrong could I be? *Is the Visual World a Grand Illusion? Spec. issue of Journal of Consciousness Studies*, 9(5-6), 17–28.
- Dennett, D. C. (2007). Philosophy as naive anthropology: Reply to Bennett and Hacker. In *Neuroscience and Philosophy: Brain, Mind, and Language* (1 edition.). Columbia University Press.
- Descartes, R. (1985b). Meditations on First Philosophy. In J. Cottingham, R. Stoothoff, & D. Murdoch (Trans.), *The Philosophical Writings of Descartes: Volume 2*. Cambridge, UK: Cambridge University Press.

- Descartes, R. (1985a). Optics. In J. Cottingham, R. Stoothoff, & D. Murdoch (Trans.), *The Philosophical Writings of Descartes: Volume 1*. Cambridge, UK: Cambridge University Press.
- Drayson, Z. (2014). The personal/subpersonal distinction. *Philosophy Compass*, 9(5), 338–346.
- Dretske, F. (1981). *Knowledge and the Flow of Information*. Cambridge, Mass.: The MIT Press.
- Dretske, F. (1995). *Naturalizing the Mind*. Cambridge, Mass.: The MIT Press.
- Dreyfus, H. L. (2013). The myth of the pervasiveness of the mental. In J. K. Schear (Ed.), *Mind, Reason and Being-in-the-World* (pp. 15–40). London: Routledge.
- Evans, G. (1982). *The Varieties of Reference*. Oxford: Oxford University Press.
- Evans, G. (1985). Molyneux’s question. In *Collected Papers*. Oxford: Oxford University Press.
- Fodor, J. A. (1983). *The Modularity of Mind: An Essay on Faculty Psychology*. Cambridge, Mass.: The MIT Press.
- Fodor, J. A. (1990). A theory of content I: The problem. In *A Theory of Content and Other Essays* (pp. 51–88). Cambridge, Mass.: The MIT Press.
- Gangopadhyay, N., & Kiverstein, J. (2009). Enactivism and the unity of perception and action. *Topoi*, 28(1), 63–73.
- Geach, P. T. (1969). *God and the Soul*. London: Routledge.
- Geldard, F. A., & Sherrick, C. E. (1972). The cutaneous “rabbit”: A perceptual illusion. *Science*, 178(4057), 178–179.
- Gendler, T. S., & Hawthorne, J. (2002). *Conceivability and Possibility*. Oxford: Oxford University Press.
- Glock, H.-J. (2008). Animal minds: Conceptual problems. Presented at the Inaugural Lecture, University of Zurich.
- Goldman, A., & de Vignemont, F. (2009). Is social cognition embodied? *Trends in cognitive sciences*, 13(4), 154–159.
- Goodman, N. (1978). *Ways of Worldmaking*. Brighton: Harvester.

- Grush, R. (2007). Time and experience. In *The Philosophy of Time*. Frankfurt: Klosterman.
- Harvey, I., Husbands, P., & Cliff, D. (1994). Seeing the light: Artificial evolution, real vision. In D. Cliff, P. Husbands, J.-A. Meyer, & S. W. Wilson (Eds.), *SAB94 Proceedings of the Third International Conference on Simulation of Adaptive Behavior: From Animals to Animats 3*. Cambridge, Mass.: The MIT Press.
- Hohwy, J. (2013). *The Predictive Mind*. Oxford: Oxford University Press.
- Hurley, S. (1998). *Consciousness in Action*. Cambridge, Mass.: Harvard University Press.
- Hurley, S., & Noë, A. (2003). Neural plasticity and consciousness. *Biology and Philosophy*, 18, 131–168.
- Hurley, S., & Noë, A. (2007). Can Hunter-Gatherers Hear Color? In M. Smith, R. Goodin, & G. Geoffrey (Eds.), *Common Minds: Essays in Honor of Philip Pettit* (pp. 55–83). Oxford: Oxford University Press.
- Husserl, E. (1991). *On the Phenomenology of the Consciousness of Internal Time (1893 - 1917)*. (J. B. Brough, Trans.). Dordrecht: Springer.
- Hutto, D. D., Kirchhoff, M., & Myin, E. (2014). *Extensive enactivism: Why keep it all in?* Unpublished manuscript.
- Hutto, D. D., & Myin, E. (2012). *Radicalizing Enactivism: Basic Minds Without Content*. Cambridge, Mass.: The MIT Press.
- Jackson, F. (1982). Epiphenomenal qualia. *Philosophical Quarterly*, 32(127), 127–136.
- James, W. (1981). *The Principles of Psychology, Volumes I and II*. Cambridge, Mass.: Harvard University Press.
- Kant, I. (1998). *Critique of Pure Reason*. (P. Guyer & A. W. Wood, Trans.). Cambridge, UK: Cambridge University Press.
- Kelly, S. (2005). The Puzzle of Temporal Experience. In *Cognition and the Brain: The Philosophy and Neuroscience Movement*. Cambridge, UK: Cambridge University Press.
- Kelly, S. D. (2008). Content and constancy: Phenomenology, psychology, and the content of perception. *Philosophy and Phenomenological Research*, 76(3), 682–690.

- Kenny, A. J. P. (1971). The homunculus fallacy. In M. Grene (Ed.), *Interpretations of Life and Mind* (pp. 65–74). London: Routledge.
- Khalidi, M. A. (2007). Innate cognitive capacities. *Mind & Language*, 22(1), 92–115.
- Kohler, I. (1962). Experiments with goggles. *Scientific American*, 206, 62–72.
- Kosslyn, S. M. (1994). *Image and Brain: The Resolution of the Imagery Debate*. Cambridge, Mass.: The MIT Press.
- Lehar, S. (1998). Gestalt isomorphism and the primacy of the subjective perceptual experience. *Behavioral and Brain Sciences*, 21(06), 763–764.
- Lehar, S. (2003). Gestalt isomorphism and the primacy of subjective conscious experience: A gestalt bubble model. *Behavioral and Brain Sciences*, 26(04), 375–408.
- Lewis, D. (1983). Postscript to “Mad pain and Martian pain.” In *Philosophical Papers: Volume 1* (pp. 130–132). Oxford: Oxford University Press.
- Lindberg, D. C. (1976). *Theories of Vision from Al-Kindi to Kepler*. Chicago: University of Chicago Press.
- Llinás, R., & Ribary, U. (1994). Coherent 40-Hz oscillation characterizes dream state in humans. *Proceedings of the National Academy of Sciences of the United States of America*, 90(5), 2078–2081.
- Lycan, W. G. (1987). *Consciousness*. Cambridge, Mass.: The MIT Press.
- Marr, D. (1982). *Vision*. Cambridge, Mass.: The MIT Press.
- Martin, M. (2008). Commentary on Action in Perception. *Philosophy and Phenomenological Research*, 76(3), 674–681.
- Matin, E. (1974). Saccadic suppression: A review and an analysis. *Psychological Bulletin*, 81(12), 899–917.
- Maturana, H., & Varela, F. (1980). *Autopoiesis and Cognition: The Realization of the Living*. Dordrecht: Springer.
- McConkie, G. W., & Zola, D. (1979). Is visual information integrated across successive fixations in reading? *Perception and Psychophysics*, 25, 221–24.

- McDowell, J. (1994a). *Mind and World*. Cambridge, Mass.: Harvard University Press.
- McDowell, J. (1994b). The content of perceptual experience. *The Philosophical Quarterly*, 44(175), 190–205.
- McGinn, C. (1991). *The Problem of Consciousness: Essays Toward a Resolution*. Oxford: Blackwell.
- Merleau-Ponty, M. (2013). *Phenomenology of Perception*. (D. A. Landes, Trans.). London: Routledge.
- Millikan, R. G. (1984). *Language, Thought, and Other Biological Categories: New Foundations for Realism*. Cambridge, Mass.: The MIT Press.
- Millikan, R. G. (1993). Compare and Contrast Dretske, Fodor, and Millikan on Teleosemantics. In *White Queen Psychology and Other Essays for Alice* (pp. 123–134). Cambridge, Mass.: The MIT Press.
- Milner, D., & Goodale, M. A. (1995). *The Visual Brain in Action*. Oxford: Oxford University Press.
- Montero, B. (1999). The body problem. *Noûs*, 33(2), 183–200.
- Müller, J. (1838). *Handbuch der Physiologie des Menschen für Vorlesungen*. Coblenz: J. Hölscher.
- Myin, E. (2001). Constrained inversions of sensations. *Philosophica*, 68(2), 31–40.
- Myin, E., & Degenaar, J. (2014). Enactive vision. In L. Shapiro (Ed.), *The Routledge Handbook of Embodied Cognition* (pp. 90–98). London: Routledge.
- Myin, E., & O'Regan, J. K. (2002). Perceptual consciousness, access to modality and skill theories: A way to naturalize phenomenology? *Journal of Consciousness Studies*, 9(1), 27–45.
- Nakayama, K. (1994). James J. Gibson - an appreciation. *Psychological Review*, 101(2), 329–335.
- Noë, A. (2002). Is the Visual World a Grand Illusion? *Is the Visual World a Grand Illusion? Special issue of Journal of Consciousness Studies*, 9(5-6), 1–12.
- Noë, A. (2004). *Action in Perception*. Cambridge, Mass.: The MIT Press.
- Noë, A. (2006). Experience of the world in time. *Analysis*, 66(289), 26–32.

- Noë, A. (2008). Reply to Campbell, Martin, and Kelly. *Philosophy and Phenomenological Research*, 76(3), 691–706.
- Noë, A. (2010). Vision without representation. In N. Gangopadhyay, M. Madary, & F. Spicer (Eds.), *Perception, action, and consciousness: Sensorimotor Dynamics and Two Visual Systems* (pp. 245–256). Oxford: Oxford University Press.
- Noë, A. (2012). *Varieties of Presence*. Cambridge, Mass.: Harvard University Press.
- Noë, A., & Thompson, E. (2004). Are there neural correlates of consciousness? *Journal of Consciousness studies*, 11(1), 3–28.
- O'Brien, G., & Opie, J. (2009). The role of representation in computation. *Cognitive Processing*, 10(1), 53–62.
- O'Regan, J. K. (1984). Retinal versus extraretinal influences in flash localization during saccadic eye movements in the presence of a visible background. *Perception & Psychophysics*, 36(1), 1–14.
- O'Regan, J. K. (1992). Solving the “real” mysteries of visual perception: The world as an outside memory. *Canadian Journal of Psychology*, 46(3), 461.
- O'Regan, J. K. (2011). *Why Red Doesn't Sound Like a Bell: Understanding the Feel of Consciousness*. Oxford: Oxford University Press.
- O'Regan, J. K., & Block, N. (2012). Discussion of J. Kevin O'Regan's “Why red doesn't sound like a bell: Understanding the feel of consciousness.” *Review of Philosophy and Psychology*, 3(1), 89–108.
- O'Regan, J. K., Myin, E., & Noë, A. (2005). Skill, corporality and alerting capacity in an account of sensory consciousness. *Progress in Brain Research*, 150, 55–68.
- O'Regan, J. K., Myin, E., & Noë, A. (2004). Towards an analytic phenomenology: The concepts of “bodiliness” and “grabbiness.” In A. Carsetti (Ed.), *Seeing, Thinking and Knowing* (Vol. 38, pp. 103–114). Springer Netherlands.
- O'Regan, J. K., & Noë, A. (2001a). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24(5), 883–917.
- O'Regan, J. K., & Noë, A. (2001b). Acting out our sensory experience. *Behavioral and Brain Sciences*, 24(5), 1011–1030.

- O'Regan, J. K., Rensink, R. A., & Clark, J. J. (1999). Change-blindness as a result of “mudsplashes.” *Nature*, *398*(6722), 34–34.
- Penrose, R. (1994). *Shadows of the Mind: A Search for the Missing Science of Consciousness*. Oxford: Oxford University Press.
- Pessoa, L., Thompson, E., & Noë, A. (1998). Finding out about filling-in: A guide to perceptual completion for visual science and the philosophy of perception. *Behavioral and Brain Sciences*, *21*(6), 723–802.
- Pettit, P. (2003). Looks as powers. *Philosophical Issues*, *13*(1), 221–252.
- Philipona, D. L., & O'Regan, J. K. (2006). Color naming, unique hues, and hue cancellation predicted from singularities in reflection properties. *Visual neuroscience*, *23*(3-4), 331–339.
- Port, R. F., & Gelder, T. V. (1995). *Mind as Motion: Explorations in the Dynamics of Cognition*. Cambridge, Mass.: The MIT Press.
- Prinz, J. J. (2006). Putting the brakes on enactive perception. *Psyche*, *12*(1), 1–19.
- Putnam, H. (1975). The meaning of “meaning.” *Minnesota Studies in the Philosophy of Science*, *7*, 131–193.
- Ramachandran, V. S., & Gregory, R. L. (1991). Perceptual filling in of artificially induced scotomas in human vision. *Nature*, *350*(6320), 699–702.
- Ramsey, W. M. (2007). *Representation Reconsidered*. Cambridge, UK: Cambridge University Press.
- Rensink, R. A., O'Regan, J. K., & Clark, J. J. (1997). To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science*, *8*(5), 368–373.
- Roberts, T. (2010). Understanding “sensorimotor understanding.” *Phenomenology and the Cognitive Sciences*, *9*(1), 101–111.
- Rorty, A. O. (1971). Not every homunculus spoils the argument. In M. Grene (Ed.), *Interpretations of Life and Mind* (pp. 75–80). London: Routledge.
- Rowlands, M. (2006). *Body Language: Representation in Action*. Cambridge, Mass.: The MIT Press.

- Rowlands, M. (2010). *The New Science of the Mind: From Extended Mind to Embodied Phenomenology*. Cambridge, Mass.: The MIT Press.
- Rupert, R. D. (2011). Embodiment, consciousness, and the massively representational mind. *Philosophical Topics*, 39(1), 99–120.
- Ryle, G. (1949). *The Concept of Mind*. London: Hutchinson's University Library.
- Ryle, G. (1979). *On Thinking*. (K. Kolenda, Ed.). Oxford: Blackwell.
- Searle, J. R. (1994). Animal Minds. *Midwest Studies In Philosophy*, 19(1), 206–219.
- Seth, A. K. (2014). A predictive processing theory of sensorimotor contingencies: Explaining the puzzle of perceptual presence and its absence in synesthesia. *Cognitive Neuroscience*, 5(2), 97–118.
- Shoemaker, S. (1980). Causality and properties. In P. Van Inwagen (Ed.), *Time and Cause* (pp. 109–135). Dordrecht: Reidel.
- Silberstein, M., & Chemero, A. (2012). Complexity and extended phenomenological-cognitive systems. *Topics in Cognitive Science*, 4(1), 35–50.
- Simons, D. J., & Chabris, C. F. (1999). Gorillas in our midst: Sustained inattention blindness for dynamic events. *Perception*, 28(9), 1059–1074.
- Sprevak, M. (2013). Fictionalism about neural representations: *The Monist*, 96(4), 539–560.
- Stanley, J., & Williamson, T. (2001). Knowing how. *Journal of Philosophy*, 98(8), 411–444.
- Summerfield, C., Trittschuh, E. H., Monti, J. M., Mesulam, M.-M., & Egner, T. (2008). Neural repetition suppression reflects fulfilled perceptual expectations. *Nature Neuroscience*, 11(9), 1004–1006.
- Tanji, J., & Evarts, E. V. (1976). Anticipatory activity of motor cortex neurons in relation to direction of an intended movement. *Journal of Neurophysiology*, 39(5), 1062–1068.
- Teller, D. Y., & Pugh, E. N. (1983). Linking propositions in color vision. In J. D. Mollon & L. T. Sharpe (Eds.), *Colour Vision: Physiology and Psychophysics*. London: Academic Press.
- Thelen, E., & Smith, L. B. (1996). *A Dynamic Systems Approach to the Development of Cognition and Action*. Cambridge, Mass.: The MIT Press.

- Thompson, E. (2007). *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*. Cambridge, Mass.: Harvard University Press.
- Tye, M. (2000). *Consciousness, Color and Content*. Cambridge, Mass.: The MIT Press.
- Van Gelder, T. (1998). The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences*, 21(5), 615–628.
- Varela, F. (1999). The Specious Present: A Neurophenomenology of Time Consciousness. In J. Petitot, F. J. Varela, B. Pacoud, & J.-M. Roy (Eds.), *Naturalizing Phenomenology* (pp. 266–314). Stanford: Stanford University Press.
- Varela, F., Thompson, E., & Rosch, E. (1991). *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge, Mass.: The MIT Press.
- Ward, D. (2009). The agent in magenta. *Psyche*, 15(1).
- Ward, D., Roberts, T., & Clark, A. (2011). Knowing what we can do: Actions, intentions, and the construction of phenomenal experience. *Synthese*, 181(3), 375–394.
- Ward, D., & Stapleton, M. (2012). Es are good: Cognition as enacted, embodied, embedded, affective and extended. In F. Paglieri (Ed.), *Consciousness in Interaction: The role of the natural and social context in shaping consciousness* (Vol. 86). Amsterdam: John Benjamins.
- Wheeler, M. (2005). *Reconstructing the Cognitive World: The Next Step*. Cambridge, Mass.: The MIT Press.
- Wheeler, M. (2010). In defence of extended functionalism. In R. Menary (Ed.), *The Extended Mind*. Cambridge, Mass.: The MIT Press.
- Wheeler, M. (2013). Science friction: Phenomenology, naturalism and cognitive science. *Royal Institute of Philosophy Supplements*, 72, 135–167.
- Wheeler, M. (2014). *The revolution will not be optimised: Enactivism, embodiment and relationality*. Presented at the 50th Anniversary Convention of the AISB, London.
- Wheeler, M., & Clark, A. (1999). Genic representation: Reconciling content and causal complexity. *British Journal for the Philosophy of Science*, 50(1), 103–135.
- Whitham, E. M., Fitzgibbon, S. P., Lewis, T. W., Pope, K. J., DeLosAngeles, D., Clark, C. R., ... Willoughby, J. O. (2011). Visual experiences during paralysis. *Frontiers in Human Neuroscience*, 5.

Wittgenstein, L. (2009). *Philosophical Investigations*. (G. E. M. Anscombe, J. Schulte, & P. M. S. Hacker, Trans.) (4th Edition edition.). Malden, Mass.: Wiley-Blackwell.

Wolfe, J. M. (1997). Visual experience: Less than you think, more than you remember. In C. Taddei-Ferretti & C. Musio (Eds.), *Neuronal Bases and Psychological Aspects of Consciousness: Proceedings of the International School of Biocybernetics* (pp. 165–185). Singapore: World Scientific.