Cognitive constraints of chimpanzees' theory of mind

Kresimir Durdevic

A thesis submitted for the degree of PhD
at the
University of St Andrews

2025

Identifier to use to cite or link to this thesis:

DOI: https://doi.org/10.17630/sta/1201

Candidate's declaration

I, Kresimir Durdevic, do hereby certify that this thesis, submitted for the degree of PhD, which is approximately 48,000 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 any degree. I confirm that any appendices included in my thesis contain only material permitted by the 'Assessment of Postgraduate Research Students' policy.

I was admitted as a research student at the University of St Andrews in September 2019.

I received funding from an organisation or institution and have acknowledged the funder(s) in the full text of my thesis.

Date Signature of candidate

12/08/2024 Kresimir Durdevic

Supervisor's declaration

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of PhD in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree. I confirm that any appendices included in the thesis contain only material permitted by the 'Assessment of Postgraduate Research Students' policy.

Date Signature of supervisor

Permission for publication

In submitting this thesis to the University of St Andrews we understand that we are 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. We also understand, unless exempt by an award of an embargo as requested below, 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 this thesis will be electronically accessible for personal or research use and that the library has the right to migrate this thesis into new electronic forms as required to ensure continued access to the thesis.

I, Kresimir Durdevic, confirm that my thesis does not contain any third-party material that requires copyright clearance.

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

Printed copy

No embargo on print copy.

Electronic copy

No embargo on electronic copy.

Date Signature of candidate

12/08/2024 Kresimir Durdevic

Date Signature of supervisor

Underpinning Research Data or Digital Outputs

Candidate's declaration

I, Kresimir Durdevic, understand that by declaring that I have original research data or digital outputs, I should make every effort in meeting the University's and research funders' requirements on the deposit and sharing of research data or research digital outputs.

Date Signature of candidate

12/08/2024 Kresimir Durdevic

Permission for publication of underpinning research data or digital outputs

We understand that for any original research data or digital outputs which are deposited, we are giving permission for them to be made available for use in accordance with the requirements of the University and research funders, for the time being in force.

We also understand that the title and the description will be published, and that the underpinning research data or digital outputs will be electronically accessible for use in accordance with the license specified at the point of deposit, unless exempt by award of an embargo as requested below.

The following is an agreed request by candidate and supervisor regarding the publication of underpinning research data or digital outputs:

No embargo on underpinning research data or digital outputs.

Date Signature of candidate

12/08/2024 Kresimir Durdevic

Date Signature of supervisor

General acknowledgement

I want to thank the Grapes lab for hosting this PhD project, and Prof Josep Call for the supervision. Also, the Edinburgh Zoo / Budongo Research Unit and Twycross Zoo teams, and the Psychology and Neuroscience workshop for their support in realising the projects. Special thanks to Kaustabh Baruah and Christina Ruiz-Mendoza for volunteering their time towards reliability coding, and to Sophie Harrower for providing cursory sanity-checks of some early drafts.

Less constrained to the present project, I owe a huge debt of gratitude to Dr Ljerka Ostojić for my methodological upbringing and induction to comparative cognition.

I was also lucky to have a group of friends who provided much emotional and intellectual support, which made the time feel more meaningful. Not exhaustively, and in alphabetical order, they are Cameron Burton, Kaustabh Baruah, Martin Schiemer, Matthew Henderson, Rasmus Overmark, Sadie Tenpas, Sophie Harrower, and Tina Petersen.

Funding

This project was funded by the ESRC, project reference 2267016 and the St Leonard's postgraduate grant, University of St Andrews.

Abstract

Like much other cognition, social cognition of great apes, or more specifically, chimpanzees (*Pan troglodytes*) embodies markers of both rich cognitive character and (somewhat) systematic limitations. This is opportune for a psychological theory that aims to outline the character of the chimpanzee's mind both in its opportunities and characteristic limitations. The theoretical space of theory of mind – the ability to represent other animals' minds – has been populated with accounts that delimit nonhuman and precocious human abilities from older humans. I deploy this heuristic in empirical (Chapters 2-4) and theoretical (Chapter 5) investigations of systematic representational limitations in chimpanzees' representational abilities. More specifically, in Chapter 2, I investigated their abilities to represent misleading appearances of objects and failed to fully replicate their reported success in tracking apparent size transformations of food items. In Chapter 3, I developed a communicative interaction task that leveraged chimpanzees' reactions to different violations of their food requests. Across two experiments, the chimpanzees failed to show a sensitivity to violations of communicative intentions that cannot be explained in instrumental reference. In Chapter 4, I adapted a different communicative task that leveraged pragmatic factors to tease out chimpanzees' tendencies to disambiguate their manual pointing gestural acts. I failed to find evidence of active disambiguation for the recipient's benefit. Finally, in Chapter 5, I developed a comprehensive and tractable system of cognitive constraints that might explain performance limitations of nonhuman primates found in the literature, as well as in the present project.

Contents

Chapter 3: Chimpanzees' sensitivity to violations in specifically communicative intentional requests 44

Chapter 1: General introduction

1.1 Explaining chimpanzees' theory of mind

A scientific explanation of a cognitive phenomenon needs a set of theoretically and operationally tractable constraints. A perceptual experience, for example, may be explained by describing its proprietary informational input, as well as the means of its transformation into psychological experience. This should also go for 'higher' cognitive processes like holding others' mental representations about the world. In a representational situation, a mental state stands in an asymmetrical relation to something in the world so that the state is about, or means something in the world. The function of representation is not exhaustively described in terms of resemblance to a causal relation, or a concurrence with its referent in the world. In order for a state to represent something, it must be somehow decoupled from what it represents, be that in the physical or mental world. What it represents, the referent, is represented *as* something, i.e., the representation allows for the situation to be conceived in a certain way: it adds information that is unique to the representational structure and cannot be reduced to its 'objective', physical characteristics (Perner, 1991).

The aim of this thesis was to contribute towards an explanation of how chimpanzees (*Pan troglodytes*) represent the content of others' minds: whether they can treat other animals as agents that hold mental states about the world. Representation of others' minds has been referred to in cognitive psychology as theory of mind (ToM), mentalising, mindreading, etc. These terms are not completely interchangeable but are also not unambiguously deliminated. They reflect different attitudes regarding how mental representation should to be instantiated psychologically. For example, 'theory of mind' sometimes implies a model akin to a scientific theory; the family of 'theory theories' liken cognition to scientific testing and updating of hypotheses (e.g., Gopnik & Wellman, 1992). 'Mindreading' may refer to a more cognitively efficient, but also a limited process (e.g., Butterfill & Apperly, 2013). Theory of mind includes different types of mental states, like beliefs, knowledge states, perceptual affordances, goals, desires, etc., which imply different representational structures and epistemic opportunities. These should be somehow constrained by their presumably species-typical psychological instantiation. In other words, if representational ability varies across species, so will the type of mental states that a species can represent. Perception, knowledge, and belief mental states are of particular importance in this introductory review, as well as further chapters, because they constitute a spectrum of representational kinds that is well placed to inform the organisation of the theory of mind problem space. One can consider how an animal sees, knows about, or believes in presumably the same referent, like a physical object, and imply related, but non-trivially different representational structures.

These differences, then, should obtain empirically as well as conceptually.

An important methodological difficulty arises in operationalising theory of mind situations. Mental states are often closely correlated with 'external', non-mental states of affairs, so that an agent can be informed of an event upon seeing it, or her information can guide her behaviour in the physical world. Crucially, however, it may not be enough to explain a mental state purely in terms of physical information like external events and actions. In so far as mental states are representational, they should be decoupled from the physicality of their referents, and include an internally specified, psychological element.

In a seminal study, Premack and Woodruff (1978) sought to examine a chimpanzee's ability to represent others' goal states. The authors showed Sarah, a trained chimpanzee, several video clips of a human performing different actions, e.g., reaching for out-of-reach food. She could then select from an array of photos showing a solution to the problem from the videos. Sarah correctly matched the photograph solutions with the appropriate videos, which alone is compatible with the idea that she was thinking of the agent's goal states. However, it is also possible that she was tracking contingencies of external events or actions – so that, keys tend to co-occur with locks, ladders with objects in high places, etc. The task, then, may have fallen short of operationally dissociating external contingencies from their internal, mental corollaries (Dennett, 1978).

One methodological solution was formulated in the so-called false belief (FB) task, where an agent is made to have an outdated belief about something in the world, like the location of an item, or the contents of a container. In a typical change-of-location FB task, initially designed for human children, a protagonist was shown to hide an item in one location, and unbeknownst to her, another agent moved the item to another location. When the protagonist returned, the child had to indicate where she thought the protagonist would search for the item. If she indicates correctly that the protagonist would search for the item in the old location, she is thought to decouple the physical, external state of affairs – where she knows or believes the item really is – from the protagonist's belief about the same state of affairs. Children exhibit a characteristic developmental milestone in which they start passing the typical false belief task between 4 and 5 (Wellman et al., 2001). Nonhuman primates, on the other hand, were reported to fail closely matched false belief tasks that children passed (Call & Tomasello, 1999), suggesting that ToM, as measured by a FB task, is a late developing, and possibly human-unique cognitive ability.

1.2 Representing perception, knowledge, and equivalent mental states

The false belief task requires the participant to represent an agent's belief state incongruent with her reality, but there is a range of other 'simpler' mental states. For one, in order to have a belief about something in the physical world, one must have some perception of it. We know that great apes and other primates understand something about how agents perceive the world: they follow other agents' gaze, observe how it may interact with the physical world, and expect it to be informative of something in the world. Great apes follow the direction of an experimenter's (Call et al., 1998; Povinelli & Eddy, 1996; Tomasello et al., 2001) or a conspecific's gaze (Okamoto-Barth et al., 2007; Tomasello et al., 1999), and do so in a way that they consider that their line of sight does not pass opaque barriers (Okamoto-Barth et al., 2007; Povinelli & Eddy, 1996; Tomasello et al., 1999, 2001). Their gaze following response is extinguished after only several demonstrations (Call et al., 1998; Povinelli & Eddy, 1996; Tomasello et al., 2001), which is compatible with them having an expectation about the gaze being informative, and when the model's gaze is shown not terminate at any (new) content, the following response is extinguished. Further, several gaze following tasks reported instances of 'double looks': after first following the gaze upwards into empty space, the apes looked back down towards the model, presumably to 'check back' after not finding anything (Bräuer et al., 2007; Call et al., 1998; Okamoto et al., 2004). Both extinction and double look rates develop ontogenetically: while infant and adult chimpanzees and rhesus monkeys alike followed gaze, infants did not extinguish their response across trials (Tomasello et al., 2001); in Brauer et al. (2005) infants were also less likely to show double looks than adults. Thus, apes read others' gaze as being directed at something in the external world, and they consider the imagined physicality of a line of gaze, so that, e.g., it cannot pass through opaque barriers.

Great apes' gaze following qualifies as a case of tracking Brentanian intentional attitudes, i.e., they relate toward things outside themselves via imagined physical relations (Crane, 1998; Gómez, 2009). By virtue of being partially imagined relations, Brentanian intentions do not need all components of an intentional attitude – the agent, the object of their intention, the direction of their gaze, etc. – to be immediately present. Instead, they can be inferred, or imagined: Maclean and Hare (2012) report how chimpanzees and bonobos can infer the object of a human actor's intention so that they will search for an alternative, novel object behind a barrier when the model exclaims in surprise towards where they saw a familiar object being baited. Presumably, they imagined an alternative, novel object based on the actor's surprised attitude. Flexibility in reading intentions is also demonstrated in non-straightforward gaze direction cues: chimpanzees begged more when an experimenter faced away from them and instead looked at them through a mirror, compared to when she faced away to a non-reflective surface. Plausibly, they computed the reflected line of sight and used it to guide their begging gestures, presumably using personal experience with mirrors (Lurz et al., 2018). Curiously, though, chimpanzees failed to include self-experience information into their gaze following response (Karg et al., 2015b): they had an experience with two handheld screens whose surface seemed identical from afar, but when examined closely, one was revealed to be opaque, while the other transparent. The screens had a differently coloured rim that helped distinguish them in the test where the experimenter gazed to the side through each screen. The chimpanzees followed the direction of her gaze as often for the transparent and opaque screens, which failed to show that their self-experience of barrier types informed their line of sight computation.

In a more dynamic, competitive setting with conspecifics, subordinate chimpanzees competed with dominant conspecifics for food items in the centre of a room. Conditions varied which items were visible to both or, or just the subordinate so that she sometimes had an ‘advantage’, i.e., a chance to steal the hidden item before the dominant. Across several experiments, the subordinates ‘stole’ more items that only they could see than the openly baited ones, suggesting that they were exploiting what their competitor could and could not see (Hare et al., 2000). In a replication attempt, however, D’Arcy and Povinelli (2002) found that the subordinates did not prefer hidden pieces of food as in the Hare et al. (2000) experiments. Bräuer et al. (2007), however, noted that the distance between food items in the replication (Karin-D’Arcy & Povinelli, 2002) was smaller than the original and applied the factor of distance to the design. At shorter distances, the subordinates secured less food, plausibly because the dominant could displace them or monopolise both items in smaller spaces. At larger distances, however, the original effect obtained. Thus, apes readily represent and exploit information about what others can and cannot see in varied and dynamic situations.

A potential ‘upper limit’ of their gaze representation may emerge in cases of reading signalers’ cues of intention to communicate. In an eye-tracking test, Kano et al. (2018) presented great apes with a video in which a model produced an attention-getting signal and gazed towards a target object on either the left or right side of the screen. The attention-getting signal was either non-social (an animated colourful shape spinning) or social (the model calling out the subject’s name). The apes’ first looks were congruent with a correct gaze following response – they looked first towards the target of the model’s gaze. They also looked equally as long to either the target or distractor objects, though. Their overall looking duration, likewise, was greater when the model produced a social cue over a non-social cue. This pattern might be due to the apes’ preference for reading information from the environment over that which is internally (psychologically) determined. Instead of preferentially attending to the object that a model is also attending to, the apes attended to all available environmental information, albeit more so in the face of a social attentional cue.

Aside from live perception, apes can also track others' off-line knowledge states, i.e., what information others have of the world when they are not immediately perceiving it. In a design similar to Hare et al. (2000), subordinates competed with competitors who saw and therefore knew about where an item was baited. They stole more food when they had privileged knowledge of the item locations compared to the dominant. They also stole more food when the dominant who saw the baiting was switched with another, ignorant competitor, compared to when they competed with an informed dominant. Interestingly, when two items were baited behind two barriers, and then one was replaced to the other without the competitor seeing it, the subordinates obtained the same amount of food as when the competitor was ignorant (Hare et al., 2001).

In a more structured competitive interaction, chimpanzees took turns choosing between baited containers against a conspecific competitor – a setup dubbed 'chimp chess'. The protocol varied whether and which locations the competitor saw, as well as which of the competitors chose first. Thus, a focal chimpanzee chose either before or after either a knowledgeable or an ignorant competitor. The focal chimpanzees avoided the location at which they both witnessed the baiting more often when they chose second than when they chose first, suggesting that they were tracking what their competitor saw and knew about the location of the food (Kaminski et al., 2008). In a similar setup with a human competitor, the food was hidden under a slanted wooden plank, and also under a second plank that laid flush against the table because the food was hidden in a secret hole under the plank. Only the subject, but not the competitor, saw this baiting, and when the subjects chose second, they showed a preference for the furtively baited location, presumably recognising that the competitor would have chosen the visibly baited (slanted) location (Schmelz et al., 2011).

These data constitute compelling evidence of great apes' abilities of representing others' perceptual and knowledge states, specifically knowledge states resulting from perceptual experience. For present purposes, perception and knowledge are an important punctuation in apes' theory of mind. Protocols testing knowledge-like state representation stop short of confronting incompatible views on reality (of the subject and the participant). Confronting incompatible mental states is the territory of belief-involving tasks. This operational progression may also be psychologically relevant – indeed, some formulation of the orthodox view divides the problem space at this point – great apes may possess theory of mind in a wider sense, including representing others' knowledge, perception, goals, etc., but not in a narrower sense of belief attribution (Call & Tomasello, 2008).

1.3 Representing beliefs

There are several relevant belief-involving tasks for great apes. Kaminski et al. (2008), in experiment 2, developed the chimp chess task to induce a false and true belief in the competitor. Both the participant and the competitors saw an initial baiting, but in two crucial conditions, in the competitor's absence, a food item was either moved to a different location (inducing a false belief in the competitor) or was moved and re-baited in the same location (inducing a true belief in the competitor). The chimpanzees did not differentiate between these conditions. Failure in the false belief condition is compatible with the orthodox view, but the true belief failure seems puzzling at first. Conventionally, true belief conditions served as controls in false belief tasks: the protagonist would witness the crucial change or transfer of an object, whereas she would miss it in the false belief condition. However, such conditions are sufficiently described in knowledge state terms: the protagonist sees, and therefore knows the last relevant change. Kaminski et al. (2008) offer a good dissociation between knowledge and true belief cases, both methodologically and empirically: the chimpanzees chose appropriately when the competitor saw the baiting, but not if the food item was removed and placed back in the competitor's absence. In the former case, a change of state happened in the competitor's absence, but even though the baiting location remained the same, the subject chimpanzees failed to exploit this information. The competitor's belief *happens to be* true, in the sense that it did not arise from having witnessed all the relevant events. Plausibly, this decoupling makes the representational structure of a justified true belief more similar to a false belief, and under the orthodox view, chimpanzees may not be well-equipped to represent it.

In line with this idea, Fabricius et al. (2010) report that children start passing the elicited response justified true belief tasks even later than the false belief tasks, although Oktay-Gür and Rakoczy (2017) report concurrent emergence when task demands are matched more closely. Justified true beliefs might be similar, or equivalent to belief states in terms of representational structure.

Relatedly, Horschler et al. (2019, 2021) examined a similar 'breakdown' across several experiments measuring looking duration as a proxy of violation of an expectation (VoE). VoE presents subjects with an event that is either in violation, or is compatible with a principle (e.g., that agents will act in accordance with their true beliefs). If participants who saw the violating event look longer at the incompatible event than those who saw the compatible event, they are thought to display a novelty response, and therefore some degree of psychological grasp of the relevant principle. Horschler et al. (2019) showed rhesus macaques a live scene where an experimenter witnessed a baiting of a piece of fruit in a location, and then, in her absence, the fruit either moved in and out of the box, or the box was moved and replaced. Across two experiments, the monkeys looked longer at a presentation where the experimenter searched at the incorrect location (violation) after presentations where the box moved, but not after the presentation where the object moved. This indicates that the monkeys formed an expectation of the agent to act congruently with her knowledge when the box moved in her absence, but not when the fruit moved. Authors interpret the latter to be a case of interrupted 'awareness relation' between an agent and the object, i.e., a break in the tracking of the relation between the agent and the object baiting.

Consider that the condition where the fruit leaves the box and the Kaminski et al. (2008) condition which replaced the food item to the same location both constitute the justified true belief case – a change of state, and a subsequent restoration of the prior state in the protagonist's absence. It is more parsimonious to interpret both under the same term of justified true belief case which causes breakdown in representation for nonhuman primates.

In a follow-up study, Horschler et al. (2021) presented two different manipulations of the target object during the agent's absence: either the object was transformed geometrically, so that an artificial flower changed shape to a blossom, or non-geometrically so that a lemon appeared to change colour by rotating rapidly. The monkeys looked longer at the violation (experimenter reaching for the wrong location) after the non-geometric manipulation but not after the geometric manipulation, suggesting that changing the object's shape, or identity (bud to blossom) was crucial to shifting the representational situation from knowledge to true belief, and likewise, causing a breakdown in the subjects' mental state tracking.

Other VoE tests also showed negative performance on tasks that initially tested infants: Marticorena et al. (2011) reported that rhesus macaques looked longer when an experimenter searched for an object in the wrong location when she had an accurate knowledge of its location, but not when she had a false belief about it. Similarly, Martin and Santos (2014) reported that a related sample of macaques failed to show a sensitivity to the difference between a violation or confirmation of a model's belief. Krachun et al. (2009) also reported great apes' failure in a competitive false belief task that 5-year-old children passed. In Krachun et al. (2010), chimpanzees failed in a change of content false belief task that 3.5-to-4.5-year-old children passed. Taken together, there is a sizable body of evidence showing that nonhuman primates consistently fail to represent other agents' beliefs, true or false, across a wide range of different protocols.

1.4 Recent data from spontaneous response methods

In contrast to these failures, modified FB protocols that measure spontaneous behavioural responses showed some positive performance in primates. The methodological advancement was made in developmental psychology where researchers measured spontaneously occurring behaviour such as duration of looking at stimuli that violate certain psychological principles (VoE), anticipatory-looking towards goals of unfinished actions, or spontaneous helping. This allowed testing preverbal infants and nonverbal animals on psychological constructs that were thus far examined in populations of older children and adult humans.

Krupenye et al. (2016) used a video-based eye tracking paradigm to show great apes false belief scenarios, and measure their anticipatory-looking patterns before agents completed their (false-belief-informed) actions. Across two experiments, the apes' first eye saccades were biased towards the false-belief-congruent location, which is compatible with them representing the actor's belief, as well as with 2-year-olds' comparable pattern in an equivalent protocol (Southgate et al., 2007; but see Kampis et al., 2020). Macaques also showed comparable performance (Hayashi et al., 2020). Heyes (2017) suggested that 'behaviour reading' – a domain-general associative mechanism – may have guided the belief-congruent anticipatory-looking patterns. For example, in experiment 1 (Krupenye et al., 2016), the apes may have tracked the concurrence of an actor's green shirt and the location of the focal object without representing any mental states. Krupenye, Kano, et al. (2017) adapted the original video stimuli to preserve only abstract and not agentive features. These videos with abstract features failed to elicit a looking pattern comparable to the previous study, suggesting that non-abstract, non-agentive stimuli were not sufficient for forming belief-compatible action predictions. Another way to control for behaviour reading is to deploy a self-experience protocol in which the participant can involve some personal experience in representing another's mental state, i.e., project her own mental state onto an agent (Heyes, 1998). Kano et al. (2019) also obtained a false-belief-congruent anticipatory-looking pattern in a version of the task which required inferring perceptual affordance based on real-life experience with types of barriers. In these videos, an actor stood behind a sheet of paper during the change of location phase of a FB video scenario. In a between-subjects design, the apes had prior experience with a similar paper sheet which was, crucially, either translucent or opaque. The apes in the opaque condition formed false-belief-congruent responses, but the translucent group did not (appropriately, since the object left the scene entirely at the end of the scene, leaving no object to approach). Thus, two anticipatory-looking studies provided evidence suggesting that behaviour reading likely was not the driver of false-belief-congruent anticipatory-looking (Krupenye et al., 2016).

Another successful methodological adaptation from developmental psychology is the application of the spontaneous helping paradigm where an experimenter elicits help from the participant in an indirect and non-specific way. Crucially, the way in which the participant chooses to help the experimenter is contingent on how she represents the experimenter's mental state. In such a task (Buttelmann et al., 2017), apes could help an experimenter in obtaining an object from inside locked containers. The experimenter hid an object in one of two boxes on the table, then either failed to witness (false belief condition) or witnessed (true belief condition) an assistant moving the object to the other box. When the experimenter then tried to open the first, now empty box, the apes helped differently between the conditions: in the false belief condition, they helped by opening the other box where the item really was – plausibly, understanding that she wanted the object, but looked in the wrong location. In the true belief condition, however, when the experimenter knew about the object's location, but tried the old location regardless, the apes were at chance in choosing which box to open. This is partially compatible with children's performance on an equivalent task: 18-month-olds helped the experimenter in the false belief condition by opening the other box, where the item was replaced, and in the true belief condition, they helped by opening the old box. 16-month-olds' data was more similar to the apes', in that they helped appropriately in the false belief condition, but were not different from chance in the true belief condition (Buttelmann et al., 2009).

Finally, the VoE method was adapted to test rhesus monkeys, but not yet great ape population, likely due to practical constraints. These studies were already reviewed, but briefly, Marticorena et al. (2011) obtained a pattern of looking duration compatible with knowledge, but not also belief representation in a design modelled after Onishi and Baillargeon (2005); Martin and Santos (2014) also failed to obtain a pattern compatible with sensitivity to others' beliefs in a design modelled after Kovács et al. (2010). Thus, VoE failed to provide evidence compatible with belief representation in rhesus monkeys.

Another recent false belief design examined the possible interference of holding in mind an agent's false belief on one's own first-order representations. The assumption is that, if animals arrive at others' mental states through simulation of their own, interference between the states might bias what they egocentrically believe or know towards what the other animal might believe. Lurz et al. (2009) applied this reasoning to a false belief task where chimpanzees watched as an experimenter saw a grape being hidden in a trough of hay, in either predetermined positions on the left or right side of a trough. Then, either in the true or false belief condition, an assistant replaced the grape to the opposite end of the trough, either with or without the experimenter witnessing it, respectively. When the chimpanzee then searched for the grape herself, the location where she could search was more frequently located in the old location's side of the trough in the false belief condition than in the true belief condition, which is compatible with the idea that an altercentric false belief influenced their own.

Before concluding this section, it is important to reflect on the replication problem that arose in the spontaneous response theory of mind literature. Some of the main original effects compatible with belief representation in human infants (Buttelmann et al., 2009; Onishi & Baillargeon, 2005; Southgate et al., 2007) failed to obtain in close replication attempts (for an overview of replications, see Kulke & Rakoczy, 2018). A special issue of Cognitive Development, gathered thus far unpublished replication attempts, thus alleviating the file drawer problem. What emerged was a checkered picture of the replicability of the original effects. Baillargeon et al. (2018) and Poulin-Dubois et al. (2018) discussed the gravity of the replication problem. The former highlighted subtle procedural 'deviations' from original protocols which may have caused the replication failures; in a reply, the latter reiterated the extent of null findings, as well as highlighted the fact that modifications in protocols, if reasonable and not a priori precluded, should not constitute 'deviations', but legitimate generalisations of the protocol. Generalisability to other theoretically relevant protocols should not be precluded by original effects without good reason, and validity of a protocol should not be contingent on the results. This is particularly important for comparative researchers who often generalise methods from developmental psychology to animal populations, which always require modifying human population protocols. Infant theory of mind possession is an open question, and this likely extends to animals, too. There are no independent replication attempts for nonhuman primate implicit ToM data, likely due to logistical constraints of availability of participants and practical constraints of using a sensitive looking behaviour recording apparatus. At present, the robustness of the findings is an open question.

Taken together, the empirical picture is conflicting. Nonhuman primates readily represent knowledge and ignorance in many contexts and fail to react appropriately in most interactions where their live partner or competitor has a false belief or a justified true belief. In some situations, though, apes show some spontaneous reactions compatible with false belief representation. These data from spontaneous tasks challenge what I referred to as the orthodox view, according to which nonhuman animals can represent some mental states, but not beliefs. A mixed empirical picture like this, then, might warrant a reinterpretation of the theoretical problem space. Indeed, the data from the last decades has motivated interesting theoretical accounts, some of which I will reflect on in the rest of this chapter.

1.5 Two-system account of theory of mind

Several theoretical accounts attempt to make room for both the orthodox human-unique view alongside data suggesting belief-compatible behavioural responses. According to an influential proposal (Apperly & Butterfill, 2009; Butterfill & Apperly, 2013), theory of mind can manifest in two ontologically and phylogenetically punctuated forms: minimal ToM, typical to nonhuman animals and younger children, and fully-blown ToM, typical to older children and adult humans. Only the latter allows for representing propositional beliefs, while the minimal form tracks nonpropositional belief-like 'registration' states. These states can produce equivalent action predictions as beliefs, but are only tracking relations between agents, objects, and locations, or successful perceptual encounters between agents and objects. Beliefs and registrations come in operational dissociation in situations where altercentric agents attend to objects which can appear under different apparent versions or aspects. If an agent is partially informed about only some aspects of an object, a subject that tracks her registrations will not be able to accommodate this aspectual informational mismatch.

For example, a fully-blown mindreader may observe Louis Lane saw Superman going into a house, and Clark Kent coming out onto the beach. Provided the full-blown mindreader knows that Louis doesn't know Clark is Superman, she will think (correctly) that Louis would not look for Superman on the beach. She knows that both Clark and Superman are aspects of the same object, and because Louis does not, she should not look for Superman when it appears as the Clark aspect. A minimal mindreader, however, will track Louis' registration of an entity coming into the house, and then coming out to the beach. To the minimal mindreader, aspectual information does not enter Louis' registration, and so will expect her to search for Superman on the beach. If we accept that a propositional belief should include aspectual information, the minimal mindreader's tracked registration qualifies, at best, as belief-like. In typical change-of-location FB tasks, both the minimal and fully-blown mechanisms make the same action prediction that the protagonist will search for the object in the old location. The full-blown mindreader will expect this because the protagonist has an outdated false belief; the minimal mindreader will expect it because that is where the last successful encounter happened.

Thus, a minimal mindreader should be able to pass a standard change-of-location false belief task. This can explain primates' success in eye-tracking and spontaneous helping tasks, but not also failures in other false belief tasks. From the proposal, it is not clear in which conditions would a minimal mindreader deploy registration tracking (as when they pass a FB task), and in which would they not (as when they fail a FB task). In other words, the lower end of minimal mindreading may be underspecified.

It is also not immediately clear how the account may accommodate nonhuman primates' justified true belief case failures (Horschler et al., 2019; Kaminski et al., 2008). If a minimal mindreader tracks agents' last valid registrations, i.e., the last location of a perceptual encounter, it would seem that it should resist 'breaking down' in the face of a change of state, like in the false belief task's change of location. If this is true, then the minimal account would fail to explain nonhuman primate failures in justified true belief failures.

While focusing on the aspectual protocol, Apperly and Butterfill's (2009; Butterfill & Apperly, 2013) proposal is not limited to the aspectual false belief task. It is also sensitive to the distinction between level 1 and 2 visual perspective taking: the former refers to cases where an animal represents *that* someone can see something, and the latter also specifies *how* someone sees something, therefore containing information specific to an agent's perspective (Flavell et al., 1983; Flavell, 1986). This distinction maps well onto the minimal and aspectual mental state tracking: a minimal mindreader tracks the fact that there is a relation between an agent and an object; the-full blown mindreader also represents how it is apprehended by the agent, i.e., what aspect of a situation is the agent aware of. The move between level 1 and 2 perspective taking to minimal and full-blown mindreading is likely just one of generalising from online visual perspective to offline epistemic states (although see Cole & Millett (2019) where the difference between perspective taking and epistemic states is contested).

Great apes show ambiguous performance in their level 2 perspective taking. Based on the fact that apes can understand that breadsticks can appear shorter or longer than they really are when partially occluded to different degrees (Karg et al., 2014), Karg et al. (2016) devised a competitive chimp-chess like perspective taking task that operationally approaches level 2 perspective taking. The focal chimpanzees had privileged perceptual access, seeing two breadsticks of the same length, whereas, for the competitor, they were partially occluded so that one looked longer than the other. Between the social and control non-social condition, the chimpanzees chose the longer-appearing stick more in the social than the non-social condition; this difference, though, was likely due to a preference for the longer-appearing stick in the non-social condition, reason for which is unclear. Regardless of their performance, apparent breadstick length need not be the best level 2 perspective taking task, as it does not require confronting two aspects (or appearances) onto one object. It seems to allow tracking which of the breadsticks a competitor might have a preference for, but not also the fact that it may appear differently to them. Just like in the false belief task, the interesting ability of level 2 perspective taking, and likewise, aspectuality, is in confronting two perspectives simultaneously.

Introducing the aspectual task effectively exacerbates the contrast in content between the egocentric and altercentric mental states and confronts them over a shared reality. This is reminiscent of the original motivation behind the false belief task, which also sought to dissociate the ego- and altercentric state contents. The aspectual protocol can be read as a more stringent version of the false belief task which 'catches' some false positives that managed to pass the task without truly confronting two incompatible mental states.

The two-system model seems to afford primates (as well as young children) a flexible mechanism for navigating the social world which in many cases resembles, but does not properly amount to, propositional belief representation. This account seems compatible with an extended orthodox view, dividing the problem space into 'proper' propositional belief and non-belief mental states. It also offers novel psychological specification of the 'minimal' psychological mechanism, which fit the latter. It may, however, stand to further specify its attitude to justified true belief cases, as well as the conditions under which registrations *fail* to deploy.

1.6 Factive and nonfactive theory of mind

An alternative way to divide the theory of mind problem space is into factive and nonfactive states (Nagel, 2017; J. Phillips et al., 2021; J. Phillips & Norby, 2019). A marker of distinction here is in how a mental state is evaluated against being true or false. In the case of knowledge, a typical factive state, false states are impossible. One cannot know something that is untrue, only be ignorant of it. Nonfactive states, however, are open to the possibility of being either true or false. They are not constrained to representing only true information; they are nonfactive, and can exist regardless of being true or false. In less language-centered terms, the means of attaining a factive knowledge state is perceptual access, whereby a subject may track an agent's perceptual affordances and whether or not they successfully relate to some factual, state of affairs. Conversely, when a subject tracks nonfactive belief states, perceptual access alone will not suffice. if one can be attributed knowledge only through 'accessing' the egocentric reality, she cannot arrive at a state incongruous with that reality. A nonfactive state must be somehow decoupled from the factual world, so that a belief can, but need not converge on it.

Such an account also seems compatible with the orthodox view of ToM possession whereby the important demarcation of theory of mind is between knowledge-like and belief-like state attribution (Call & Tomasello, 2008). Like the two-system account (Apperly & Butterfill, 2009; Butterfill & Apperly, 2013), it offers further specification for the nonhuman-typical ToM. Phillips and Norby (2019) stress two key features that any type of theory of mind must satisfy: tracking and separation. Tracking entails maintenance of another agent's mental state and readily deploying it for purposes of predicting their actions, exploiting vantage points, etc. Separation entails maintaining the differences in informational states of one's own and an agent's mental states so that, e.g., when an agent is ignorant of something that the subject is knowledgeable of, the subject cannot mistake the information states. In articulating the mechanism, the authors draw a parallel with maps: each subject's informational state is represented analogously to a private map which is kept independent of other maps. To sum, a ToM-ready animal must first track another agent's perceptual access to their own map. Next, in order to make the separation between one's own and another's mental content, one must generate another map for the agent, and remove the missing information that they do not have access to. In this way, Phillips and Norby (2021) grant the factive mindreader an altercentric representation that is still, strictly speaking, not in conflict with one's own (and therefore not false) representation of the world, while still being functionally separate from the egocentric image of the world.

Nonhuman primates and younger children should be constrained to factive ToM, i.e., knowledge-like representation, while only older children and adult humans should readily represent nonfactive states like beliefs.

Focusing on nonhuman primates, this model fails to accommodate success in spontaneous false belief tasks, but it is compatible with negative false belief and justified true belief performance. If belief states are the limiting case, a primate should not have the provision to 'persist' with a more basic form of ToM (like registration tracking) through a false belief task. This seems like a more clear-cut condition for success/failure: if a condition can be construed in terms of knowledge or ignorance, factive ToM is enough; nonfactive ToM will only be deployed for false belief and justified true belief cases.

Thus, the two-system account and the factive-nonfactive account converge on the matter of belief representation, which is in reach of only the more mature form of ToM exhibited by older children and adult humans. They diverge on the specifics of how and when are more basic forms of mental state representation deployed – either through tracking perceptual access to factual reality, or successful perceptual encounters with objects. Operationally, the minimal mindreader should be able to track an agent's registration of an object's location in a false belief task, but a factive mindreader should not. Similarly, the minimal mindreader may persist in the tracking of an agent's mental state in a justified true belief condition, whereas a factive mindreader's tracking should break down.

1.7 Other accounts

These are not the only theoretical attempts at reconciling precocious and nonhuman positive performance with the apparent performance limitations in theory of mind. Tomasello (2018), for example, offers a reading of spontaneous response false belief tasks by which such protocols do not require the participant (infant, or ape) to represent a truly false belief, one that is incompatible with objective reality. Instead, what the implicit tasks might be detecting is a representation of others' intentional attitudes. The main difference between an intentional attitude and a belief, for Tomasello (2018), is in how perspectives are coordinated: younger children and apes can take others' perspectives, or mental states, but do not also confront them. The difference being, only the latter entails simultaneously apprehending two states which can be incompatible with one another. This contrast was made evident in appearance-reality discrimination and perspective taking tasks, where in the former, 3-year-olds could recognise that objects can be 'real' or 'deceptive' in their appearance, but not also contrast two aspects of an object onto one physical object (Moll & Tomasello, 2012); in the latter, they know which items adults refer to, even when some are distorted in appearance by a colour filter, but they cannot reliably decide on how a distorted object looks like for someone else (Moll et al., 2013; Moll & Meltzoff, 2011a).

Understanding others' intentional attitudes, thus construed, seems roughly compatible with the minimal mindreading factive mindreading: limited to tracking relations within a factual reality, without the possibility of emerging conflict of perspectival information.

This is not an exhaustive review of all data or theoretical accounts pertaining to great ape theory of mind, but a selective review of that which allows for dissociation between its ‘more basic’ and ‘more complex’ forms. Similar ‘cuts’ in the theory of mind problem space are made by different authors, which speaks to the usefulness of the presumed underlying heuristic. In this thesis, I took this heuristic of dichotomising the problem space roughly between belief-like and knowledge-like states to be the guiding heuristic on the operational and theoretical levels. I developed the heuristic theoretically in more general and more specific terms and applied it to novel, and hopefully interesting, operational situations.

1.8 Summary of the coming content

In the coming chapters, I will report on three empirical studies which applied the theoretical heuristic primed in the introduction. In Chapter 2, I reported on an appearance-reality discrimination paradigm. Appearance-reality discrimination is a key ingredient of the representational structure involved in the notion of aspectuality: the ability to represent an object as having different incompatible appearances. The goal was to replicate the basic AR test of Krachun et al. (2009; Krachun et al., 2016), with the ultimate aim of developing the paradigm to inform chimpanzees’ representational constraints, hopefully amounting to some of the above discussed distinctions.

In Chapter 3, I attempted a different approach: I leveraged pragmatic factors of object choice interactions by violating specifically communicative intentions of the chimpanzees’ requests for specific food items. The crucial methodological moment was in distinguishing between the chimpanzees’ instrumental and specifically communicative intentions, thus approaching the dichotomy of external and internal reference in communicative acts.

In Chapter 4, I attempt another communicative task, an adaptation of Tauzin et al. (2020) which offered apes a chance to modify their pointing gestures in order to disambiguate their external referent for the benefit of an experimenter. This would show a sensitivity to the fact that their interlocutors can have mistaken readings of their gestural acts, thus approaching a sensitivity to internal specification of communicative intention.

Finally, in Chapter 5, I reflect more thoroughly on some of the theoretical considerations primed in the introduction with the aim of articulating a domain-general, species-specific set of cognitive constraints for chimpanzees. I also attempt to apply this core dichotomous heuristic to domains of thought outside theory of mind.

Chapter 2: Appearance-reality discrimination

2.1 Introduction

Discerning the appearance and reality of objects in the world is a developing ability in human children: Flavell et al. (1983) showed that preschool children develop an appreciation that an ambiguous object can appear as one thing, but really be another. For example, a piece of sponge painted grey might appear as a rock, but inspection reveals its true nature. 3-year-olds were largely insensitive to these distinctions, but 4- and 5-year-olds exhibited fewer errors when asked about objects' true and appearing properties or identities. Connections between performance in such appearance-reality (AR) discrimination tasks and false belief representation were apparent from the outset (Flavell, 1986); the incompatibility of an object's appearing and true properties resembles a situation where two people's perspectives clash on how they believe the world to be. Indeed, the performance of children four and older who succeed in AR discrimination tasks correlates with their performance on theory of mind tasks (including the false belief task) (Gopnik & Astington, 1988; Moore et al., 1990). This suggests that the developmental shift of AR discrimination is about confronting incompatible aspects of an object. However, AR discrimination ontogeny might be constituted by a more gradual representational change. Moll and Tomasello (2012) show that AR tasks can be passed by younger children when the task does not require confronting reality and appearance, but only knowing which objects are deceptive. Three-year-olds were shown how a chocolate bar was a true bar, whereas a deceptive object was an eraser that only looked like a chocolate bar. Then, facing the objects side by side, the experimenter asked for 'the real, true' object, and 'what only looks like' an object. Three-year-olds were above chance here when they had to discern which items were deceptive, and which were true. In the second experiment, a different sample of three-year-olds was shown an ambiguous object in the centre, and on each side, unambiguous objects corresponding to its two possible identities. They were asked to indicate which of the lateral objects is the ambiguous object like, thus having to confront two aspects onto the same, ambiguous object. This task was too difficult for the children, presumably because, unlike in experiment 1, they had to confront two incompatible aspects onto a single object. Thus, in older children, a crucial representational shift may emerge by which they are able to confront incompatible aspects of objects.

By the same token, 3-year-old children can take the perspective of another agent in a perspective taking task, but not also confront incompatible perspectives. For example, when only one of two toys is placed behind a yellow colour filter from the agent's perspective, and she ambiguously asks for 'the green one', 3-year-olds know that she wants the distorted object (Moll et al., 2013; Moll & Meltzoff, 2011a). In a 'perspective confronting' version task, however, only 4.5-year-olds succeeded in coordinating incompatible perspectives on the *same* object (Moll et al., 2013).

The limiting case of aspectual representation contrast is also well illustrated in the dual-naming task: the experimenter and the child play a game of alternate naming of an object so that one has to name an object without using a name that was already used. If an experimenter calls a plush rabbit toy 'rabbit', the child can use 'bunny' or 'hare' etc. In the control condition, they take turns naming the colours of an object, so that a flower can be green and red. The control condition does not present a true case of incompatible reference because the colours refer to adjacent, but not mutually incompatible aspects of the same object; in the experimental condition, however, each name describes the object in its whole, which renders the names incompatible. Four-year-olds can master the dual-naming task, and their performance correlated with that of a false belief task (Doherty & Perner, 1998).

Thus, AR discrimination may be supported by a similar representational structure that is responsible for the successful representation of others' beliefs. Apperly and Robinson (2003) showed that 5-to-6-year-old children find it easier to reason about false beliefs compared to when agents had true beliefs about alternative descriptions of an object. Meaning that the difficulty arose from having to track a relation between the agent and alternative descriptions, or aspects, of an object. Rakoczy et al. (2015), further, showed that children's performance in the classical false belief task correlates with their performance on a simplified aspectual task, solidifying the link between representing beliefs and incompatible aspects of entities in the world.

If representing others' beliefs is like discriminating appearances from reality or dual identities of objects, should nonhuman primates' theory of mind limitations (Call & Tomasello, 2008) predict their limitations on AR discrimination tasks? Some authors suggested aspectual theory of mind tasks as operational limits of nonhuman (or infant) theory of mind (Butterfill & Apperly, 2013; Lurz, 2009). In this formulation, nonhuman primates should readily represent relations between agents and things in the world, but not also when their relations involve objects that appear under different aspects. If an animal is limited in her theory of mind abilities, she should also be limited in her AR discrimination abilities.

Before turning to nonhuman primates' AR discrimination, I will review briefly what we know about their object individuation since having a concept of an object in the first place is necessary to qualify for representing objects' appearances as decoupled from reality. The object concept can be broken down at least into spatiotemporal, featural, and identity aspects (Wiggins, 1997). Spatiotemporality, which is likely the minimal sense in which an object concept is maintained, implies tracking objects as bounded, coherent, and three-dimensional. Minimally, then, they retain their shape and separation from the environment as they move through space, but will be relatively insensitive to how the objects look like. Featural individuation, additionally, imbues objects with descriptors like colour and shape. Finally, identity-based individuation tracks the internal properties of objects, so that they persist even as the external appearances change, for example, if a slice of banana is painted orange to look like a carrot, it is still ‘truly’ a banana. It is only on this level of object representation that the notion of AR discrimination starts being relevant because it requires decoupling objective, physical properties of an object – what it looks like – from what it really is, independently of its physical appearance.

Violation of expectation paradigms revealed that even 10-month-old human children individuate objects spatiotemporally, and by 12 months, featurally (Xu & Carey, 1996). For nonhuman primates, there is evidence that nonhuman primates use both spatiotemporal and featural information to guide their search in manual object individuation paradigms. The basic logic of these tasks works so that the information presented in the baiting, e.g., the number of items, is either congruent or incongruent with what the animal finds in a manual search. When the search does not match the baiting, then the animal may search for an object she infers should also be present. Great apes, capuchins, and rhesus macaques were shown to search compatibly with tracking spatiotemporal and featural information about objects (Flombaum et al., 2004; Kersken et al., 2020; Mendes et al., 2008, 2011; Santos et al., 2002). Rhesus monkeys also showed a looking time pattern compatible with featural individuation in an eye-tracking paradigm (Uller et al., 1997).

In order to arrive at identity-level individuation, one must either conflate internal and external information or two incompatible external aspects. Phillips et al. (2007) showed that macaque monkeys can keep track of the internal properties of objects, i.e., they matched the taste of a food piece to a fruit it came from. In a further study, they showed the monkeys a surface transformation of the fruit so that it appeared as another, inserted an artificial piece from the fruit in the box, and let the monkeys search inside. They searched for longer when the found piece did not match the fruit *under* the transformed property, compared to when the fruit and the piece in the box did match. Presumably, the monkeys were tracking whether the piece in the box matched the fruit that was transformed to look like another fruit, i.e., its internal property against its external appearance. It is possible, however, that the covering of a fruit with the shell of another fruit might not count as transformation, but occlusion, which would either make the shell easier to ignore, or it fail to impose the pertinent hierarchical structure.

In a possibly more stringent test, Cacchione et al. (2016) report on an object transformation task in which food items were painted to look like a different item. They presented great apes with a pair of food items – a carrot slice and a banana slice and transformed one of the items (by painting and covering it with a fake peel) in each condition so that the final choice was between two *apparently* identical items, one being the transformed one. Despite initial low performance, certain modifications helped the apes perform better: in one condition, increasing the item sizes, which increased the cost and benefits (and thus motivation) behind their choices, and in another, only a paper object was painted to resemble a banana, which lowered memory demands. Apes' performance was above chance in these conditions (but not in two other conditions). While the latter condition with reduced memory demands does not necessarily require contrasting incompatible aspects – the apes can pass by tracking the location of one, untransformed object – the former condition with size manipulations does require tracking transformed and true properties. Featural transformation situations coerce a new representational structure that accommodates internal identities relatively decoupled from varying appearances. In a similar transformation, Kanzi, a bonobo trained on lexigramic communication, viewed objects being transformed and indicated which of the objects corresponded to a lexigram of, e.g., either a banana or a carrot. Kanzi correctly indicated the identity of transformed objects when he saw the transformations, but not in the controls where transformation was hidden, thus outperforming the Cacchione et al. (2016) sample (Lurz et al., 2022). Notably, the Cacchione et al. (2016) study was vulnerable to a strategy by which the apes could fixate on one (presumably, pre-transformation banana) item's location and choose accordingly. In Kanzi's case, he only knew which item he was meant to indicate after the baiting, making it unlikely that he was tracking only one location. Still, both protocols are vulnerable to a spatiotemporal tracking strategy by which one could track the location of the pre-transformation states.

In a more exhaustively controlled protocol, Krachun et al. (2009) showed chimpanzees big and small grapes baited in boxes with a minimising, or a maximising lens, respectively. The more desirable big grape, then, appeared as small in a minimising lens box, and a small grape appeared large in a maximising lens box. The apes witnessed the baiting and would request either grape with a manual point. A series of test progressions controlled for alternative strategies: In the initial basic test, the two grapes were simply baited inside the boxes and offered for choice. Next, the seen and unseen tracking tests controlled for spatiotemporal tracking: in the seen tracking test, after the baiting, the boxes were stacked on each other and either switched positions or returned to the same position. This prepared the chimpanzees for the subsequent unseen tracking test, which controlled for spatiotemporal tracking by occluding the switch/return phase so that the chimpanzees could not trace the motion of the boxes. The chimpanzees, as a group, did not pass the basic test, and their performance was distributed bi-modally across sessions, by which some chimpanzees tended to consistently do well, while some tended to do poorly. Given this heterogeneity, the authors analysed the behaviour individually, weighing the performance of individuals against chance for each session. 8/14 chimpanzees passed the basic test, a further 5 passed the unseen tracking test, and 4 of those passed a further transfer task where food types were changed, which required a degree of generalisation of the lens manipulation to novel stimuli.

A further condition tested for a strategy by which the chimpanzees potentially reversed the contingency of choice and reward, so that choosing a more palatable option led to obtaining the less palatable – in this situation, picking the smaller-appearing grape to obtain the latter by simply inversing the outcome associated with the choice. This is an a priori unlikely strategy, as chimpanzees reportedly fail to choose a smaller food quantity over a larger one when the task reversed the reward contingency (Boysen et al., 2001; Boysen & Berntson, 1995). Likewise, in Krachun et al. (2009), none of the chimpanzees passed the two reverse reward contingency tests, showing that they were likely not inverting the 'natural' choice-reward contingency.

Thus, given that only a minority of the participants passed all testing stages, the chimpanzees did not show strong evidence of AR discrimination on the lens AR test. Krachun et al. (2016) extended this test to a different sample of chimpanzees, included two different types of manipulations, and extended the stopping rule to allow more exposure to the manipulations. Like in the original task, the chimpanzees had to pass the 'simpler' test stage in order to proceed further. The original design allowed only two sessions per test, with either 12 for the basic test or 6 (for the tracking tests) trials per session. The subsequent study increased the session number to 8 12-trial sessions for the basic test and 4 12-trial sessions for each tracking test. An additional avoid-lens test was introduced which controlled for a perceptually-based rule of avoiding the magnifying lens; a large grape was baited behind a magnifying lens so that it looked bigger, and a medium-sized grape was baited behind a transparent non-distorting lens. If the chimpanzees used and persisted with the strategy of avoiding a minimising lens, they would have failed in the avoid lens test. 4 12-trial sessions were allowed for this test.

All 5 of the tested chimpanzees passed all administered test phases, suggesting that AR discrimination is within their abilities given enough relevant experience. The authors included two additional manipulation types: the mirror test, in which the number of grapes was deceptively increased in a mirror box, and a colour filter test, in which a coloured screen distorted the colour of boxes containing food rewards. 1/6 chimpanzees passed all the trials in the mirror test, and 2/7 passed the colour filter tests. The lens tests with an extended stopping rule, then, showed to be an opportune test of AR discrimination which still allowed controlling for alternative strategies like spatiotemporal or other 'low level' strategies.

Therefore, I aimed to replicate the test with a new sample of chimpanzees housed in the RZSS Edinburgh Zoo in order to gauge their AR discrimination abilities, and with the distal aim of adapting the task to be sensitive to distinctions between taking and confronting perspectives, or incompatible aspects of physical objects (Doherty & Perner, 1998; Moll & Tomasello, 2012; Tomasello, 2018).

2.2 Methods

Participants

I tested 11 chimpanzees (*Pan troglodytes*) (4 female) out of 16 present in the group. Average age of the sample was 31.82, ranging between 8 and 46 years. Participation was voluntary and opportunistic, i.e., based on the willingness and ability of the chimpanzees to approach and follow through with the tests. Edith's testing was interrupted due to practical constraints, and David passed away during testing, so their data is incomplete.

Subjects were housed at the Budongo Research Unit (BRU) in RZSS Edinburgh Zoo. Testing was conducted in a designated research room connected to indoor day areas, which were connected to outdoor and bed areas. There were no restrictions to food or water access outside of the usual diet. At the beginning of each testing session, the keepers would open the door to the testing room and sound a signal to notify the chimpanzees. The whole troop (with the exception of individuals separated for welfare purposes) could access the testing room during the duration of the testing session (4 hours). There was no coercion to participate by the experimenter (me) or the keepers, except for calling the chimpanzees' names and showing them food rewards. Occasionally, keepers would distract individuals to avoid overcrowding or monopolising the space, which could lead to altercation, or hinder testing.

Materials

The basic manipulation of the test used magnifying and minimising lenses to alter the apparent size of grapes. In all but the 'mixed' sessions, the grapes were baited so that the truly bigger grape appeared smaller behind the minimising lens, and the truly smaller grape appeared bigger behind the magnifying lens. I bought grapes from the supermarket on the day prior to testing, and on the morning of the testing, I sorted them into big and small sizes. Grapes around 15mm (+/-2mm) in diameter or length were sorted as small, and those around 35 (+/- 2mm) as big. I measured a handful of prototypical grapes with a small measuring calliper and matched the remaining grapes by visual inspection. I squished each grape flat slightly so they would not roll inside the lens boxes.

The lens boxes which held the distorting lenses were 3D-printed in the St Andrews Psychology and Neuroscience workshop. The white plastic filament was made into boxes with a rectangular base and one vertical face slanted at a 30° angle. The boxes held either a minimising or a magnifying lens which could be slotted in or out. The base of the box was 170*160mm, and it was 80mm high. The lens, slotted to fit at the centre of the slanted face, had a 50mm diameter. I wiped the boxes down regularly and they were visually indistinguishable when viewed from the front and empty. A mark was made on the back to distinguish them from my point of view and the camera's which was positioned behind and over my shoulder.

Figure 2.1. Photo of lens boxes and their manipulation of two identical model objects.

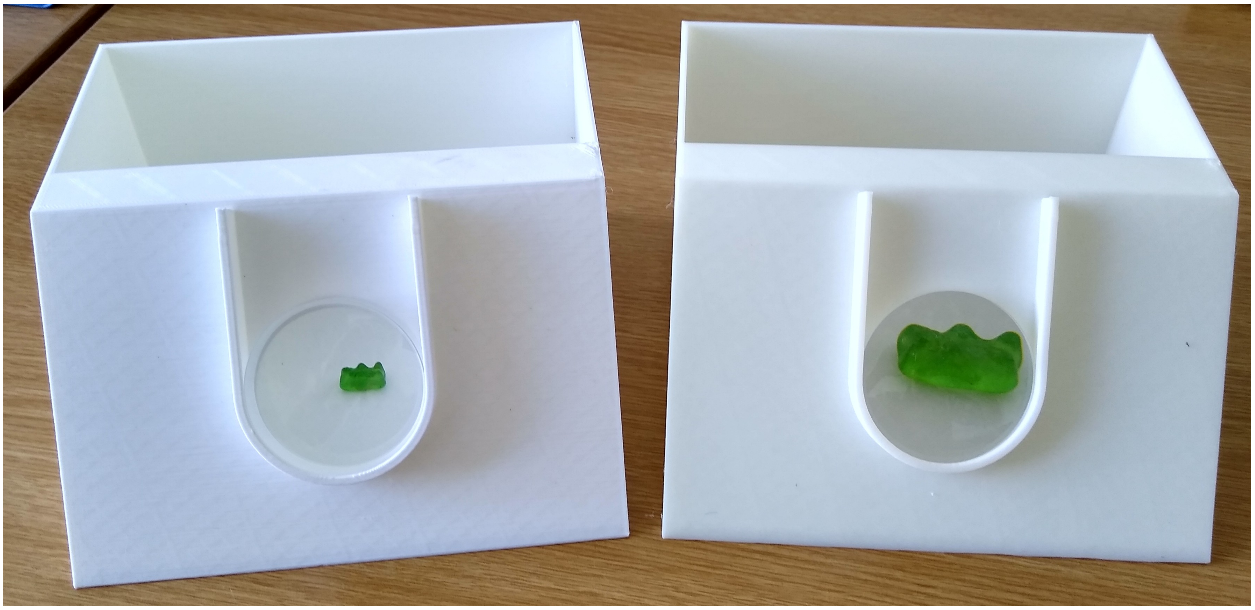

Note. The grapes baited in the actual test were big and small, not identical as the candy show above.

Procedure

The design had each participant progress through up to 4 different test session types sequentially. Each participant would move on to the subsequent stage once the criterion performance was reached within the stopping rule. Criteria and stopping rule were based on Krachun et al. (2016) and are outlined in table 2.1. If a participant failed to reach criterion within the allotted number of sessions, they did not proceed with further sessions, and wold be given routine cup game interactions as enrichment instead.

Table 2.1. Number of sessions allowed to reach criterion and the passing criterion per session type.

Session type	**Sessions before stopping**	**Passing criterion**
Food Preference	Until impractical	2 consecutive 10/12 correct
Basic	8	2 consecutive 9/12 correct or 1 10/12 correct
Seen tracking	4	2 consecutive 9/12 correct or 1 10/12 correct
Unseen tracking	4	2 consecutive 9/12 correct or 1 10/12 correct
Unseen tracking mixed	2	NA

Note. Each chimpanzee progressed through different session types (shown here in descending order), provided she reached the criterion performance within the stopping rule predetermined for each session type.

Before each test trial of any type (1-4), I administered two preference trials and two demonstration trials. Preference trials served to determine the motivation of the chimpanzees; if they did not choose the big grape in the first two trials, two more trials were presented, and I would require 3/4 big grape choices before allowing the chimpanzee to proceed to test trials. If they did not reach 3/4 correct trials, I stopped the session and engaged in enrichment cup game interactions until the next testing day. In the subsequent demonstration trials, I baited each lens box slowly, lifting the grapes out and back in to the box twice before leaving them inside. I then moved on to the test session regardless of their choices in the demo trials. Moving on to the test sessions was not contingent on performance in demonstration trials.

0. Preference: I placed a large and a small grape on the left or right positions on the sliding panel, and the chimpanzees chose between the pair. In baiting, I varied the order in which I baited the left and right sides pseudo-randomly so that no more than 3 consecutive trials would have the same side baited first.

1. Basic test: I placed the lens boxes on the left and right sides of the sliding panel, and then baited the small grape under the magnifying lens box, and the big grape under the minimising lens box using plastic tongs. The left-right positions of the boxes and the order of baiting each side were counterbalanced pseudo-randomly so that the same configuration (with respect to all dimensions) would not appear in more than 3 consecutive trials. Passing this stage served to make sure that the chimpanzees tracked the task dynamic; other low-level explanations are plausible at this stage, e.g., spatiotemporal tracking. Further tests served to disambiguate between possible low-level explanations.

2. Seen tracking: the trial proceeded as the basic trial, but after the baiting, I stacked the baited lens

boxes on top of each other, the right always on top of the left. Then, I either switched the boxes to an opposite position, or returned them to the same position (switch/replace). The point of these trials was to prepare the chimpanzees for the subsequent phase where the switch is occluded – they would come to expect that stacking event leads to a switch in position in 50% of the trials.

3. Unseen tracking: the trial proceeded as in the seen tracking test, but after stacking the boxes, I raised an opaque occluder to hide the switch / replace stage. After I unstacked the boxes, I removed the occluder and pushed the sliding panel to allow a choice. Spatiotemporal tracking could not lead to the correct choice in these trials.

4. Unseen tracking mixed: these trials proceeded just as in the unseen tracking conditions, but in half of the trials, the baiting contingency was reversed from big grape-minimising lens and small grape-maximising lens to big grape-maximising lens and small grape-minimising lens. In these 'reversed' trials, the big grape appeared bigger, and the small grape appeared smaller. In the other half, the standard trials, the baiting proceeded as in prior trials. This way, had the chimpanzees learned to choose the grape that appeared smaller, or tracked some other perceptual information correlating with the minimising lens, they would have chosen the wrong grape in the trials with the reversed configuration. The standard and reversed baiting configurations were interspersed across two sessions, 6 trials of either kind in each session, so I collected 12 trials of each type across two sessions and weighed their performance of 12 trials within a trial type.

2.3 Results

I used R version 4.4.2 (2024-10-31) for all analyses.

A separate coder re-coded 15% of the videos; inter-observer reliability was high, $\chi^2 = 0.88$, $p < 0.05$.

Like in Krachun et al. (2016), I analysed the performance of each chimpanzee individually. Testing proceeded sequentially, i.e., different tests entailed different 'stages' of the task's complexity; a participant had a predetermined number of sessions in which she could 'pass' and progress onto the subsequent stage. If she had reached criterion within the allotted number of sessions, she would not proceed with further testing.

Passing criteria was determined as an above chance number of 'correct' trials in a binomial test at the $p < 0.05$ criterion, i.e., either 10/12 correct trials within one session or across two subsequent sessions with 9/12 correct trials.

11 chimpanzees showed a preference for the large grape in the pretests and proceeded to the first, basic test. 11/11 chimpanzees passed the basic test, and 10/11 passed the seen tracking test (Figure 2.2). 6/11 chimpanzees passed the unseen tracking test, progressing to the mixed trials. No chimpanzees passed both the standard and reversed trials, suggesting that they relied on some perceptual factor in making AR discriminations in the lens tests (Figure 2.3). David's testing stopped due to his death; Edith's testing was stopped due to practical constraints of testing session allowance.

Figure 2.2. Number of correct trials broken down per subject and session type.

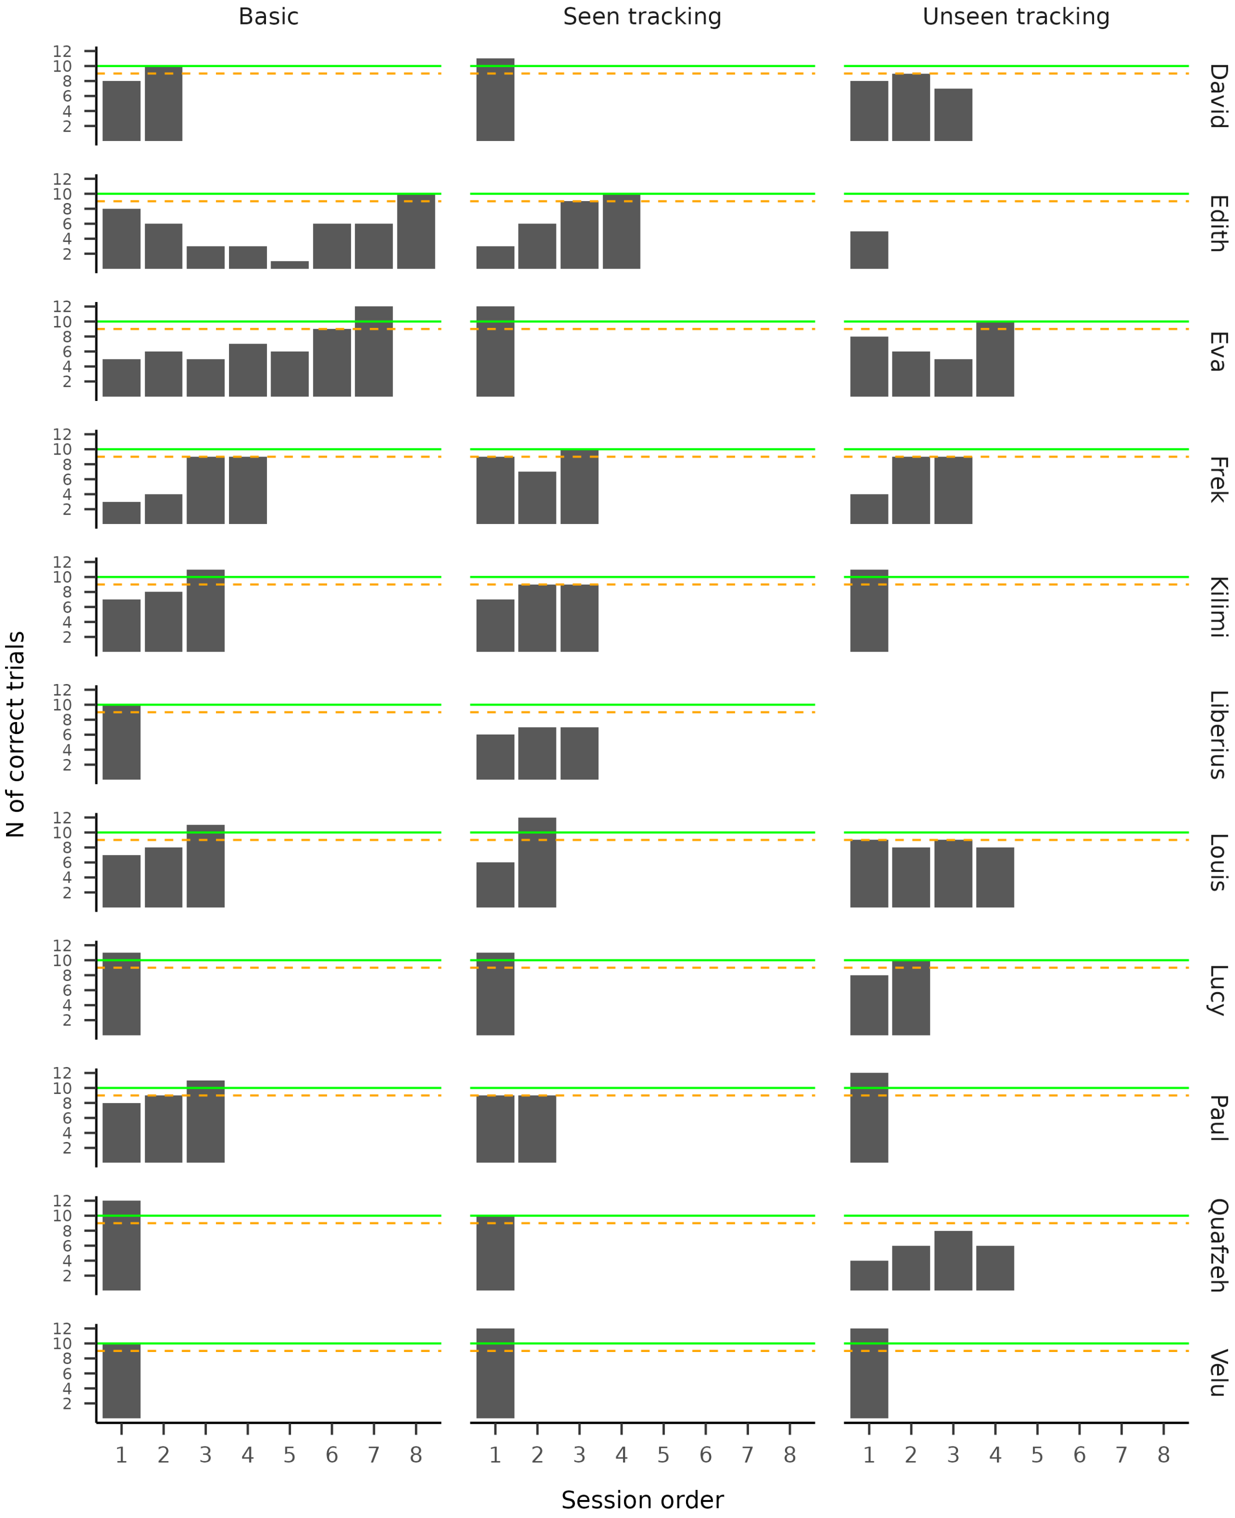

Note. This figure shows the frequency of correct trials (chosen truly large grape) broken down by subject and session type. The two horizontal lines designate the criterion for progressing through the sessions: either two subsequent sessions with at least 9/12 correct trials (dashed orange line), or one session with 10/12 correct trials (green full line). Edith's testing was interrupted because it was impractical to continue data collection for her alone; David passed away during testing; Liberius' fourth session is not shown because testing access was paused in the middle of his fourth session, but it had too many wrong trials to reach criterion and was not picked up when testing restarted.

From inspecting Figure 2.2, it is apparent that there is much inter-individual heterogeneity in the number of trials chimpanzees took to reach criterion. Velu, for example, was at the far end of good performance, reaching criterion in first session of each test type. At the other end, e.g., Edith's performance in the basic test resembled an inverted U-curve, and improved gradually in the seen tracking test.

Regarding performance in mixed trials specifically, only Eva passed the reversed trials of mixed sessions, but not standard trials of mixed sessions; the former require flexibly adjusting the strategy from previous sessions, whereas the latter do not, which casts doubt on whether Eva really understood the nature of the manipulation. 4/6 chimpanzees that made it to the mixed trials was successful in the standard trials, but not the reversed trials, which is compatible with the idea that they were using a 'low level' perceptual strategy, instead of flexibly adjusting to the manipulation. None but Eva were successful in the reversed trials.

Figure 2.3. Performance of the 6/11 chimpanzees in the mixed trials (standard and reversed).

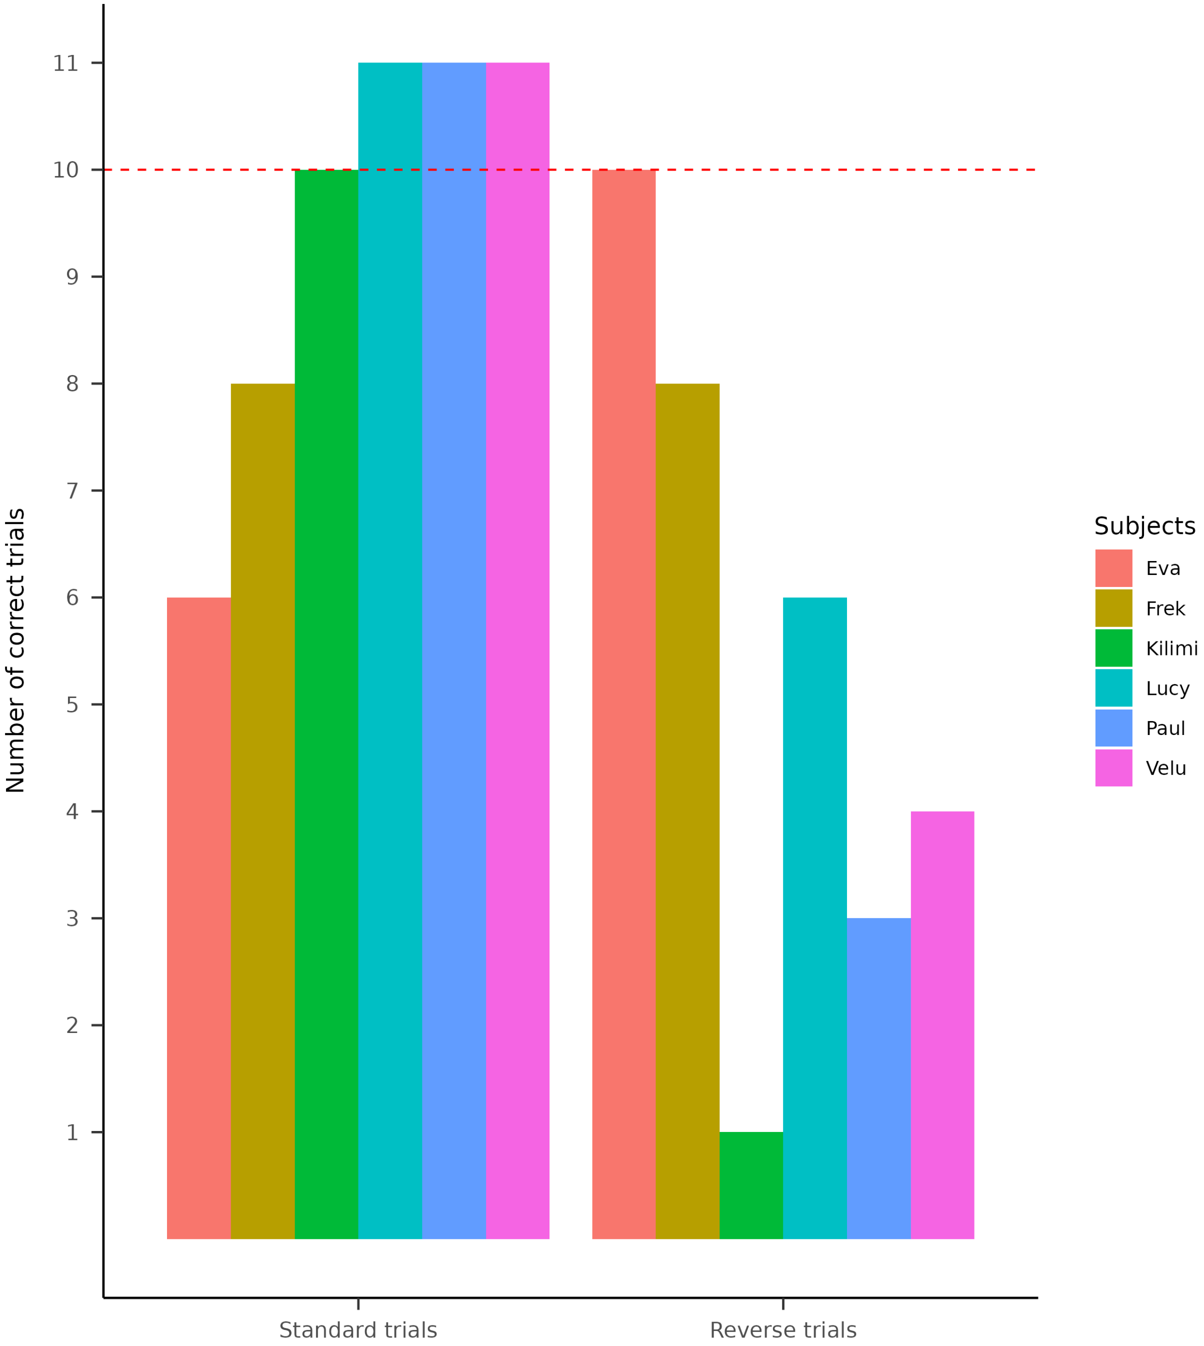

Note. Standard trials had the usual baiting contingency: small grape - maximising lens, big grape - minimising lens. Reverse trials reversed the contingency: small grape - minimising lens, big grape - maximising. Thus, in the latter, the small grapes looked smaller, and big grapes, bigger. Only two sessions of mixed trials were administered, for a total of 12 trials per type.

2.4 Discussion

Overall, the performance on the magnifying-minimising lens task is likely to be driven by tracking lower-level perceptual cues. None of the chimpanzees passed all of the relevant perceptual control trials, which is compatible with the notion that they were using some perceptual cues to guide their choices instead of considering how each grape was distorted anew in each trial. All (11/11) chimpanzees passed the basic test, which only shows a general understanding of the task dynamic, and is relatively ambiguous as to the deployed cognitive or perceptual strategy. Only 6/11 chimpanzees passed the unseen tracking tests, which precluded the use of spatiotemporal tracking.

In the mixed trials, I tried to disambiguate as to whether the 6/11 chimpanzees that passed the unseen tracking tests used a deflationary strategy such as always choosing the grape that appears smaller behind the lens, or by tracking some factor that correlated with the lens affordances. 3/6 chimpanzees passed the standard trials, which were identical to previous unseen tracking tests, but interspersed across two sessions with reverse trails, which reversed the baiting contingency. In reverse trials, the smaller and bigger grapes were transformed to be even smaller and bigger, thus the chimpanzees could not have chosen the smaller appearing grape (or tracked a correlated perceptual factor). Only Eva passed the reverse trials, but given that she did not pass the standard trials also, it is likely that her performance may have been due to chance.

The chimpanzees' failure at the reverse trials need not imply that they were using a rule such as the 'avoid small image' strategy. Plausibly, other perceptual cues also co-varied with the presentation of a grape behind a minimising lens. These factors include subtle differences in the lenses' manipulations – e.g., light refraction, viewing angle affordance, differences in visual space distortion, etc. As one of the less obvious possibilities, consider the differing viewing angles afforded by the lenses. Compared to the minimising lens, the maximising lens has a narrower range of angles from which its distorted object is visible. This means that, as the angle from which the lens is looked at increases (e.g., by moving one's head to the side), the object behind the minimising lens will disappear from view sooner than the one behind the maximising lens. This is not a problem for a modest viewing angle from which most chimpanzees likely viewed the lens boxes most of the time, but it is possible that this led to more perceptual exposure to the grape behind the minimising lens, which was the correct choice for all but the final reverse trials. This is just one possible perceptual cue that co-varied with the lens manipulations and may not have been successfully controlled for.

Another alternative strategy considered already in the original Krachun et al. (2009) paper is the reverse reward contingency strategy by which the typical contingency a choice of an object and the delivery are reversed, and a subject chooses the inverse of the desired option. This is not a highly plausible candidate for a psychological explanation, however: chimpanzees were reported to struggle with reverse reward contingency in quantity judgment tasks (Boysen et al., 2001; Boysen & Berntson, 1995), and in the lens test context (Krachun et al., 2009).

Had some of the chimpanzees passed the mixed trials, it would not necessarily mean that they would have been using a targeted strategy of flexibly confronting appearing and real aspects of an object. Moll and Tomasello (2012) discuss a strategy by which, instead of contrasting two aspects of an object, AR discrimination is possible by tracking which of the objects is the 'real' one and which is 'fake', i.e., marking them as such, and not contrasting incompatible real and appearing aspects of objects. Such marking entails a relatively sophisticated object representational strategy, but it does not amount to tracking incompatible aspects of objects, which was the target phenomenon presently. In the current lens task, one could track the big or small grapes as they enter the lens boxes and now appear to be of inverse sizes, and attach 'big' and 'small' labels to the distorted objects. This type of representation goes beyond the immediate or 'face value', but it need not create an aspectual situation in the narrower, more cognitively sophisticated sense.

The lens task is not the only available nonverbal AR task. Krachun et al. (2016) used two additional manipulations: mirrors distorting the apparent object number, and colour filters distorting objects' apparent colour. These manipulations, however, are less compelling cases of aspectual representational structure. When one considers the number of items, one need not entertain two versions of an object, but add or subtract objects that may be there. The colour filter may also fail to amount to the same structure. It need not change the identity of an object, only distort the retinal image of everything behind a coloured screen. The 'illusion' is not transparent in its effect: it need not 'fool' the observer into believing that the object is of the same colour as the colour filter, only that there is some kind of colour distortion at hand. If I see a yellow object behind a green screen, it will not make me believe it is green, likely I will perceive it as a stained yellow. The space around the object will be tinted too, betraying what colour is being imposed by the filter. Effectively, the filter will mark the object as having a misleading appearance, but it will not also transform its identity. With this in mind, a colour filter manipulation may not be aspectual, but might instead require 'marking' objects as misleading, in line with the less sophisticated strategy referenced in Moll and Tomasello (2012).

Further, Karg et al. (2014) used a test with partially occluded breadsticks of different lengths; when placed behind an opaque barrier, from the subject's point of view, the shorter bread stick stuck from outwith the barrier more than the longer one, thus appearing longer. This manipulation, however, is vulnerable to interpretations of spatiotemporal tracking.

Transformation tasks like the Cacchione et al. (2016) also offer a good testing avenue but do not control for alternative strategies like spatiotemporal tracking. The items are presented to the subjects, and one is painted to resemble the other. Since the objects do not change positions, tracking the location of the initially desirable item is sufficient for successful discrimination.

The lens test might remain the best option for approaching aspectuality with an AR discrimination protocol. Unlike the mirror manipulation, the lens modifies an object and creates its alternative appearing counterpart; unlike the colour filter task, its transformation does not betray its nature, so it would not seem to encourage the strategy of 'marking' the objects as misleading and true. Finally, unlike transformation through painting or covering, it affords spatiotemporal tracking control. Great apes also perform better in this manipulation than others (Cacchione et al., 2016; Krachun et al., 2016). In so far as the lens manipulation is opaque, and the minimised grapes really look like a small grape, the manipulation entails a transformation, and may approach an aspectual situation implied in Apperly and Butterfill's (2009; Butterfill & Apperly, 2013) notion of aspectuality (but in an egocentric situation).

There were two possible plans to take the AR design further: a) to develop a lens manipulation equivalent to Moll and Tomasello's (2012) task which distinguishes between a strategy of marking misleading or true objects and one that contrasts different aspects of an object; and b) to apply the lens test to a social setting so that it meets the operational criteria laid out by Apperly and Butterfill (2009; Butterfill & Apperly, 2013). This was contingent on positive (and compelling) AR discrimination evidence, as per data reported by Krachun et al. (2016). Performance of the present sample – with high individual variation, and with no chimpanzees passing all perceptual controls – and the difficulty of controlling all relevant perceptual factors count as reasons against using a lens-based AR discrimination protocol to inform more nuanced theoretical questions.

Chapter 3: Chimpanzees' sensitivity to violations in specifically communicative intentional requests

3.1 Introduction

The matter of internal reference, or internal mental representation, is situated at an important fulcrum of debates in comparative cognition: to what extent do nonhuman animals hold internal, mental, and relatively disembodied representations of things in the world? The question, at least when asked within the loose bounds of classical cognitivism, usually strings the theoretical space between two poles: on one hand, there is a tendency to explain behaviour in terms of associative mechanisms that react to or use stimuli from the external world without also having internal mental representations of external events (e.g., Penn et al., 2008); on the other, one invokes at least some degree of disembodied mental representation which cannot straightforwardly be explained in terms of its corresponding external stimuli. Both sides of the contention – the 'lean' associative on one hand, and 'rich' cognitivist on the other – invoke arguments of parsimony with the success that is apparent when assuming their respective theoretical commitments. For the minimalist, it is parsimonious to assume that associativist processes like reward contingencies and associating concurrent cues in the environment are sufficient explanations until proven otherwise; the cognitivist approach insists that something like an internal (cognitive) representational structure is needed to guide the otherwise intractable experience and operations of the mind. The result is a theoretical impasse – can one explain nonhuman behaviour exhaustively with reference to external, environmental information, or is it necessary to constrain mental experience with some type of internal, cognitive structure?

Manual pointing gestures in communicative settings here might offer some good ground for informing the matter. In pointing, one refers to something in the external world but also purports to affect the mental state of the recipient. It pertains, then, to both internal and external content at the same time. Trained captive great apes routinely point for the benefit of experimenters in object choice interactions: typically, an experimenter presents two or more objects, or containers hiding objects, and the ape may point to one of the two, to which the experimenter responds by sharing the content that she chose. It is not immediately clear what may be the supporting psychological instantiation that guides such an interaction. To what, and how does the ape attend to in a cup game? In a particular reading, intentional communicative acts are done for the benefit of communicative recipients and serve to convey the communicator's intention towards something in the world; which will then affect the recipient to act in accordance with her goal (Wilson & Sperber, 2012). Applied to the cup game situation, a chimpanzee would (1) hold an intention towards the food item (or whatever is hiding it), and (2) hold a goal corresponding to that item. She would (3) reveal her intention to the experimenter, who will interpret the message and help her in achieving the goal. I understand intention, here, as a mental attitude that an organism holds towards something outside itself (Gómez 2009, 2008). This analysis is relatively rich with cognitive content since it requires not only (1) a first-order intention of the producer but also a (3) second-order intention whereby a chimpanzee reveals her (1) first-order intention to the experimenter. This is not the only psychological explanation, though. The chimpanzee may be persisting in an action that was reinforced with a rewarding outcome, which is compatible with an associativist explanation. Or, plausibly, they might be (1) intending towards an object, (2) desire it, but instead of (3) revealing their intention to an experimenter, they perform an instrumental action which yields a (2) desirable outcome.

It might be helpful to consider some details of the presumably richer model. Wilson and Sperber (2012) leverage the signaler's intention towards the recipient (3), or her mental state. They refer to this as second-order intentionality – the intention to communicate. This is the load bearer of human-typical intentional communication. Specifically, they argue that a speaker's (1) first-order intention towards something in the world is neither sufficient nor necessary for successful communication. If that were the case, any act of intending towards the world would count towards a communicative intention. If a conspecific orients rapidly towards a source of a loud noise, that can be informative of the state of affairs in the world, or their mental state; however, this is not because the agent communicated anything; her intention is merely a part of the world. One can exploit others' first-order intentions, but this is not sufficient for qualifying as intentional communication.

Identifying a signaler's intention is also not a necessary case for intentional communication. In some communicative cases, a recipient can recognise that a signaler is trying to communicate *something*, fails to understand the message, but can still either infer her meaning or find the 'incomplete' message informative in itself. For example, if I am travelling in a country whose language I do not speak, and a local talks to me sternly, I may try to infer based on the context that I may be doing *something* culturally inappropriate, even if I do not know what. In another context, a colleague in a meeting might struggle for the correct word while articulating a concept; even without the appropriate word, I can infer an approximation of the correct idea based on the rest of the discussion. In both cases, the content of the intention was opaque or underspecified, but contextual inference could guide me to the correct interpretation. The crucial moment, however, is in having recognised that there is something that is trying to be communicated – an (3) intention to communicate. This is the key of the ostensive inferential model of communication (henceforth, inferential model): the recipient must recognise a signaler's intention to communicate in an ostensive act, which is a trigger for the recipient to infer the communicative meaning found ‘inside’ the speaker’s mind.

To loop back into the chimpanzee cup game example situation, we may conjecture that it is not parsimonious to suppose that a chimpanzee is intending communicatively as per the inferential model. A manual point is likely intentional in a first-order, Brentanian sense, but not also communicatively intentional in the inferential (ostensive) sense. It may be informative, but it need not be specifically communicatively meaningful. It is at the level of (3) second-order intention, or ostension, that one finds specifically communicative meaning. By separating the (3) specifically communicative intention from the (1) first-order intention, we achieve a degree of separation between (1) *what* the signaler is communicating about from (3) the fact that she is intending to communicate; in other words, her specifically communicative intention is relatively independent of reference to external content. It is disembodied, it need not carry information about the world – that is the job of (1) first-order intention, or further contextual inference or guesswork.

It may seem farfetched to posit disembodied conceptual structure to be the driver of meaning in a communicative message, given that we so routinely communicate about things in the physical world. However, disembodied cognitive structures not only conform to some common sense intuitions about communication; they also converge with the fact that communicative expressions often do not refer to physical entities. Humans can think and talk about idealised geometrical shapes and other mathematical entities, or ‘abstract’ linguistic examples. These can correspond to some things in the external world – a cube may be an idealisation of a dice, a normal distribution represents a measurable trait in a population, and a word can denote a physical entity. However, crucially, one can apprehend these without physical, observational reference. In this case, physical referents are incidental. One may be more inclined to accept this for mathematical than for linguistic ideas. We routinely use language for describing the world, but ‘scientific’ statements like ‘All swans are white’ do not have a clear observational correspondence – one cannot ever round up *all* the swans, neither practically nor in principle. Likewise, one cannot corroborate a statement whose evidence is not available, like ‘Cesar crossed the Rubicon’ – past events cannot be brought back for inquiry. Psychological constructs, which we are happy to talk about as real and non-trivial, are indirectly induced from behavioural proxies. Even if these cases are in the minority, the fact that propositions can express these nonexistent referents casually should be instructive of their nature. Specifically communicative intentions benefit from this ‘inductive gap’ as it becomes easier to imagine how intention can approach disembodiedness.

Intentions that behaviourally present as acts of manual point and gaze, however, need not be equivalent to propositional structures in their disembodiedness. Manual points, as well as gazes, are always in a physical relation to their referent: an agent looks at, or points to, a target, and must adhere to some (imagined) physical relations – at the very least, e.g., whether the imaginary line of gaze intersects the target unobstructedly. Since agents hold intentions and express them in bodily cues, the intentional relation should contain information pertaining to their perception and attention to the relevant aspects of the world.

Several studies show that apes' begging gestures vary along with humans' physical cues of attentional states like their body or face orientation, or whether or not their sight is obstructed. Apes generally show more manual and visible gestures when the interlocutor is attentive or can see (Call & Tomasello, 1994; Hostetter et al., 2001, 2006; Liebal et al., 2004; Poss et al., 2006; Tomasello et al., 1994). Kaminski et al. (2004) also showed that, while apes were sensitive to the orientation of an experimenter's face, they only showed this sensitivity when their posture matched the face orientation, suggesting that body and face orientation are represented relatively independently. Bulloch et al. (2008), similarly, found that chimpanzees were sensitive to whether or not an experimenter's sight was occluded, but struggled when her posture did not match her face orientation. Great apes will also move to accommodate an experimenter's body orientation if she turns away with food: chimpanzees (but not also other apes) followed the experimenter's body orientation even when the food did not move (Liebal et al., 2004). Thus, apes track the agents' attentional and perceptual states, which shows representing the agent's end of intentional relation, i.e., that the relation contains an agent with some perceptual and attentional properties.

At the other end of the intentional relation, one must track the physicalities of intentional acts like points and gazes. At the very least, that means holding an expectation that a point or a gaze is directed at something in the world (and is not just a directional cue), and how it might intersect opaque barriers. Chimpanzees follow a human experimenter's gaze direction (Call et al., 1998; Povinelli & Eddy, 1996; Tomasello et al., 2001), as well as that of a conspecific (Okamoto-Barth et al., 2007; Tomasello et al., 1999) by orienting or searching towards where an agent is looking at. They also consider whether the imaginary line of gaze intersects opaque barriers (Hare et al., 2001; Okamoto-Barth et al., 2007; Povinelli & Eddy, 1996; Tomasello et al., 1999, 2001) as well as that gaze can be reflected in mirrors (Lurz et al., 2018). This shows consideration of the physicality of an imaginary line of gaze, which determines the content of an intention behind it.

It is also important that a gaze, and therefore an intentional relation, should be directed at a specific target. Gaze following data also shows that apes form expectations about gazes being directed at informative targets. Across different tests, the gaze following response is extinguished across trials, suggesting a habituation-like pattern (Call et al., 1998; Povinelli & Eddy, 1996; Tomasello et al., 2001). The habituated response is unlikely to be a simple non-specific co-orienting response, though: while older primates in Tomasello et al. (2001) extinguished their response across trials, the infants continued to follow gaze across trials. It could be that gaze following follows an ontogenetic trajectory, moving from a non-specific co-orienting response, and towards an understanding of intention with an expectation of some content determination. If adult primates had an expectation about gaze terminating at a meaningful target, then a gaze that fails to do that will become uninformative and uninteresting. Apes also exhibit ‘double looks’ whereby, after following a gaze and not finding anything at its endpoint, they look back to the experimenter, presumably to check back with the source of the gaze (Bräuer et al., 2005; Call et al., 1998; Okamoto-Barth et al., 2007). Bräuer et al. (2005) reported that adults were more likely to show double looks than infants, further supporting that gaze following becomes ‘more specific’ through ontogeny. Finally, Maclean and Hare (2012) showed that apes can infer the target of a gaze based on pragmatic factors. A human model was shown to be either familiar or unfamiliar with an object which was visibly baited behind one of two occluded baiting locations. When the model looked ambiguously towards the locations and exclaimed excitedly, the apes were more likely to search behind the second location if the model was familiar with the object behind the first location than when she was not. Presumably, they inferred that the excited cry was a surprise reaction to a novel object, which they searched for in the second location. Primate gaze following, then, is intentional on ‘both ends’: it concerns perceptual and attentional affordances on the agent’s end, and it encounters something in the physical world.

Apes also draw similar intentional relations in communicative interactions. They make gestural adjustments to varying spatial arrangements of referents (Gonseth et al., 2017; Tauzin et al., 2020) and in order to guide the experimenter to their preferred location (Roberts et al., 2014). This is compatible with their gestural production adhering to first-order intentionality. It is not clear, though, that apes can also read second-order, specifically communicative intentions.

Object choice tasks have shown some relevant performance limitations in apes. Typically, an experimenter presents the ape with a pair of opaque containers or occluders and baits a food item under one of them (baiting is hidden). Before offering the ape a choice between locations, the experiment provides a cue informative of the food item’s location. Call et al. (1998) showed that chimpanzees did not use the experimenter’s gaze directed at the baited location when the food was baited in cups, but did so when the food was baited behind opaque barriers, or horizontal cylinders with a closed-end facing the chimpanzees, and the open end facing the experimenter. In contrast, Itakura et al. (1998) found that chimpanzees could use various gazing and pointing cues, which together show somewhat mixed results. In further investigation of object choice cues, both Itakura et al. (1999) and Call et al. (2000) found that both local enhancement cues like approaching and standing physically next to the baited location, or various vocalisations or noises, were used more readily to identify the baited location. Notably, Itakura et al. (1999) tested mostly human-reared chimpanzees in a setup with either human or trained conspecific models; the chimpanzees were successful in using local enhancement cues, but not gaze or point cues, regardless of the model.

These data can be interpreted to suggest that, despite reliably following gaze apes do not reliably read gazes as cues of second-order, ostensive cues indicating the experimenter's intention to inform them. In contrast, they exploit some cues of local enhancement and reaching. Plausibly, the cues that apes use are those of first-order intention (in case of reaching), and non-social enhancement of locations or objects. It remains possible, however, that these limitations are task-specific: Mulcahy and Call (2009) found that increasing the distance from the model and the baiting locations, so that the experimenter points distally and not proximally, improves apes' performance on the object choice task.

In conjunction, reliable intentional (in the embodied, Brentanian sense) gaze following and unreliable specifically communicative intention reading might indicate a disjunction in great apes' communicative representational abilities. When reading gaze, one need only apply a first-order intentional attitude, i.e., that an agent is looking at something. In contrast, an object choice situation involves not only the model's first-order intention but also her intending to communicate, i.e., her second-order ostensive intention.

It could be that nonhuman primates do not take ostensive cues to be disembodied prompts for inferring communicative meaning, but prefer to interpret signals as direct evidence of what they stand in physical relation to. Does the reason for this shortcoming lie in the cognitive abilities, however, or in the extraneous constraints of manual acts? Perhaps verbal language is more suited for disembodied intentionality without requiring a unique cognitive underpinning. It would be helpful to find cases where manual intentional acts are not in direct physical relation to their intended referents in the world. Gómez' (2008, 2009) reading of Brentanian intentionality may be helpful here. For him, a case of straightforward Brentanian intentionality includes animals referring to something in the external world, but one can extend the concept to a more general category of intentional availability, by which one can intend toward things outside immediate or even imagined physical relations. There are two kinds of intentional availability: intentional inexistence and intentional nonexistence. Intentional inexistence includes 'straightforward' cases where intentional relations are drawn between immediately present, or imagined, presently absent physical entities. The latter allows for representing cases of absent reference in which one can refer to something that is currently not there but normally is there – for example, I can look at my empty coffee cup and think that I ran out of coffee, or I can observe that I depleted the ink from my pen. In both cases, I am referring to something absent from the situation. This type of reference is of the same type as that which relates to physically present entities: in either case, there is an imagined physical relation. Apes showed that they can manually point towards absent entities, i.e., towards a depleted baiting location that previously held food. Crucially, they did so more often when choosing between two different food types than between two locations of the same type (Bohn et al., 2015). The overall rate of pointing to an absent entity was relatively low, but regardless, pointing to an emptied-out plate suggests a reference to an absent entity, and therefore, supports the notion that apes' intentional communication exhausts the case of intentional inexistence.

In the second type of intentional availability, i.e., intentional nonexistence, the referent is not absent, but it positively does not exist. The difference lies in the fact that, when intending to absent entities, there is still an imagined relation to the physical reference, just one that is presently absent. In the case of intentional nonexistence, there should be no such reference. If I think of an empty coffee cup, I am still relating to coffee, which just happens to be absent. If I lack this contextual or normative information (i.e., that coffee is normally in the cup), I could say that there is nothing in the cup, whereby I am not using an absent entity to support the representation. For example, if I look at a nondescript empty box, and say 'There is nothing in this box', I lack a reference point for what is usually there; I need to think of it as truly empty. If the box had a label like the Smarties logo, or if it were a milk carton, I would have a point of reference for an inexistent intention towards an absent entity (candy or milk, respectively). However, a nondescript box does not allow abductive inference as to the likely contents, so I am likely to think of it as containing nothing. To my knowledge, there is no pointing behaviour protocol that could show intentional nonexistence suited for nonhuman animals. The Bohn et al. (2015) protocol does not go further from absent inexistence. Indeed, it is difficult to think of how a point or a gaze might persist without an imagined physical relation, or whether that is possible at all. Gómez (2008) prefers to use children's pretend play behaviour as an example that is not in reference to a physical state of affairs – children's pretend play need not physically match the real act it is purporting to emulate. Before giving up on nonexistent manual gestures, though, it may be worthwhile to attempt to explore cognitive models with equivalent representational structure and through those vindicate the use of manual gestures as useful experimental indices of disembodied intention.

Since pointing, and potentially, gazing, are communicative acts, we may be tempted to read Brentanian intentional acts in reference to the ostensive inferential communication model (Wilson & Sperber, 2012). To reiterate, under the inferential model, when a speaker intends to reveal her intention to the audience, the audience may register the direct evidence of ostension, and infer the correct information based on the context, including, but not limited to, the evidence from the utterance (first-order intention). The crucial aspect of the model, though, is the disembodied intention to communicate, i.e., the ostensive second-order intention. Simple Brentanian intentionality is equivalent to the first-order intention of the inferential model where a producer is straightforwardly intending to something present or momentarily absent from the world. The latter, in the inferential model, is indirect reference, e.g., 'it rained last Christmas' – the event is not presently available but is still accommodated in the imagined relation. Indirect reference and intention towards absent entities can both account for holding intentions to an empty food plate that normally holds desirable food pellets. Inferential intentionality, however, can amount to more radical independence from external content whereby it can be recognised by speakers as contained without specifying any content at all, not even indirectly. Nonexistence could be a representational equivalent to second-order intention as it persists without the support of a present or absent imagined physical relation.

Manual gestures and gaze can be proxies of either first or second-order intentions. For example, if I am trying to convey a message to a friend furtively, or in a very loud room, I can seek out their eye contact out and gaze somewhere specific, hoping to provoke them into attending to what I am looking at. Eye contact, here, is an ostensive cue of my intention to communicate. Such an interaction maps onto the situation in a computerised eye-tracking gaze following a task (Kano et al., 2018). The authors showed great apes a recording of a model who stood centrally behind two identical objects. The model first displayed an attention-grabbing cue: in the ostensive experimental condition, the model looked at the camera and called the participant ape's name; in the non-ostensive control, he took a bite out of an apple, followed by a loud crunching sound. The model, then, looked towards one of the objects (left or right), after which the apes' looking behaviour towards the screen was recorded. An ostensive cue served as direct evidence of the communicative intention, which might, under the inferential communication model, prompt the apes to infer the meaning of the model's communicative intent. In this case, the authors expected the ostensive cue to facilitate following the model's gaze to the target compared to the non-ostensive control. Experiment 3 revealed that the apes' first looks were more often directed at the target over the distractor, i.e., their first looks were compatible with a gaze following response. Looking duration, however, revealed overall more time looking at either target in the ostensive over the control condition, but not also a bias towards the target over the distractor. In other words, the ostensive cue caused more time looking, or inspecting, but not towards the target vs. the distractor. The authors interpret this to suggest that '...apes understand, at least partly, that humans' ostensive signals precede referential information, as they searched the environment longer after witnessing the ostensive signals compared to an equally attention-grabbing non-ostensive signals.' (Kano et al., 2018, p. 724). The subjects, thus, plausibly recognised the model's intention to communicate but failed to infer the correct referent, thus showing a limited sensitivity to ostension.

The authors further drew an analogy between this gaze bias and the emulation social learning model: when observing others using tools, apes attend more to the environmental affordances instead of specific techniques that the models are using. This is referred to as a bias for emulation over imitation: rather than entering a social learning situation with the assumption that the demonstrator has a technique in their knowledge repertoire which they will demonstrate, and an ape can imitate, they reinvent the technique that is plausible given some external state of affairs, like objects left over from tool use (Tennie et al., 2006). Emulation, here construed, could reflect a bias towards seeking information in the environment rather than ‘in the mind’ of a demonstrator. Translating this to the case of ostensive communication, the apes might have interpreted the ostensive cue not as a prompt to infer the content from the model’s mind, but to seek the environment for a plausible referent. The results of Kano et al. (2018), therefore, might reveal a bias in how apes consider ostensive information: not as a prompt to infer what an agent means, but to search the environment for plausible candidates. In most cases, these two heuristics will converge on the same outcome; in the Kano et al. (2018) task, the model gazed at featurally indistinct objects. Since the objects are featurally equivalent, the environment affords no information, and the only source of information that can distinguish them is contained in the model’s internal communicative intention. If they are unable to represent such communicative intention, they have no grounds to ‘pick out’ or infer either of the featurally indistinguishable objects as the correct referent.

If the apes really are biased towards seeking information in the environment rather than internally, they plausibly would have distinguished between two featurally distinct candidate referents. In fact, in the counterpart of this study with human infants, the testing situation differed from the Kano et al. (2018) data so that the target and distractor objects looked distinct. Here, infants followed the gaze more accurately on both first look and looking duration measurements (Senju & Csibra, 2008). Unfortunately, just as the great ape study did not include a condition with different objects, the infant study did not include a condition with identical objects, which hinders interpretations of internal vs. external communicative inference. We do not know whether infants were successful because they read ostension better than apes, or because for them, the environment contained enough information for discrimination. One might expect an animal equipped with an understanding of internal, communicative intentionality to give primacy to internal information and to successfully disambiguate between either featurally indistinct or distinct objects. If an organism is biased, and limited to, externally determined reference, then she would fail to distinguish between featurally indistinct, but succeed between featurally distinct referents.

In line with this reasoning, a 2*2 design would complete the picture between the Kano et al. (2018) and Senju and Csibra (2008) studies: in a similar gaze following task, one should vary object feature distinction (same / different objects) and ostension (ostensive / non-ostensive communicative cue). Say that an animal that can represent ostensive communicative intention is guided by the ‘inferential model’. Such an animal would identify the correct referent regardless of whether the model is looking at one of the same objects or one of different objects, and only following an ostensive cue. If an animal is limited to seeking information in the environment over in the intention of an actor, say that she is acting according to the ‘alternative model’. This animal would expect successful gaze following in the different, but not in the same condition, as environmental information is available only in the former.

Table 3.1. Conceptual design for disambiguating between the inferential and alternative communicative models.

Intentionality model	**Same objects**	**Different objects**
Inferential model	Y	Y
Alternative model	N	Y

Note. An inferential model will disambiguate between the referents in the same and different object conditions, whereas the alternative model will fail to disambiguate the referents in the same, but not the different object condition. Successful disambiguation is marked 'Y' and a failure, 'N'. Both models would expect a more pronounced gaze following response following an ostensive cue, compared to a non-ostensive cue (this factor is not shown in the table for simplicity).

If the contrast between the internal and external bias of communicative intention of the inferential and alternative models is a legitimate principle, it should also generalise to situations where apes produce the communicative act. Suppose there are two identical Mars bars on the table, and the signaler asks the recipient for 'the Mars bar', pointing to the one on the left. The recipient acknowledges her request and shares the one on the right (identical item). Imagine the signaler protests that her request is not addressed properly. The protest is sensible from the inferential communicative model – both objects carry the same 'external' affordances for the signaler (they look and taste the same), and the distinction which makes one of them 'correct' is made only when the signaler's utterance is recognised as relevant independently of the goal (obtaining a Mars bar). The signaler is right to protest receiving the 'wrong' Mars bar because the recipient is addressing only her overall goal or first-order intention, interpreting the act as ('she wants a Mars bar'), but ignores the fact that her intention is also determined in her mind (i.e., the left item). The second-order intention, in a sense overrides, the recipient's first-order goal. Under the alternative model, the signaler has a goal towards a Mars bar; the recipient then searches the space for a referent that meets this criteria and finds two. Regardless of which one she shares, the speaker has no grounds to protest as her first-order goal of obtaining a Mars bar is obtained; the fact that her point singled out one is incidental. There is no second-order communicative intention to impose a conflict.

One problem for this conceptual design is that a speaker might provoke protest by sharing the opposite item only by virtue of not matching, e.g., the spatial configuration of the point's referent. The speaker might not be considering both referents when providing her point, focusing on just the one object. When she fails to obtain it, she might protest simply for her 'narrow imperative' request not being met. Under this narrow imperative explanation, the speaker protests, not because of a violation of her second-order intention, but because she failed to obtain the one object she held in mind. Some developmental data suggests that young children may not default to this: in Grosse et al. (2010), 18-30-month-olds requested certain toys over others, to which the experimenter would respond by first verbally addressing their request, and then sharing one of the toys. What differed between the conditions was whether the verbal address and the shared toy were correct or not. In the successful interaction baseline condition, the experimenter responded to the children's request for toy A with a verbal confirmation i.e., 'Oh you want A', and handed her the correct toy, A. In the 'happy accident' condition, the experimenter pretended to misunderstand the request, saying 'Oh, you want B?' but still sharing the correct toy, A. Thus, the first-order goal (obtaining A) was satisfied in both cases, but the second-order communicative intention was misaddressed in the latter condition only. Children corrected the experimenter more in the condition where they were misaddressed than in the baseline, showing that, independently of fulfilling the goal of their interaction, they protested the experimenter's incorrect addressing of their request. Thus, even when the children's goal was fulfilled, they still protested when their communicative intention alone was misaddressed.

Protest, or gestural repair in the face of (specifically) communicative violations might be a feasible proxy of sensitivity to communicative intentionality. Rate of behaviour subsequent to a specific violation in an interaction was previously used to gauge sensitivity to purposeful actions of a cooperative experimenter: during a series of routine food delivery trials, the experimenter failed to share food in a trial, either acting unwilling (e.g., by withholding the item), or unable (e.g., by clumsily dropping the item). Chimpanzees exhibited more behaviours subsequently to unwilling trials compared to unable trials, showing a sensitivity to unexpected violations in interactions (Call et al., 2004). In this situation, though, the chimpanzees did not choose the items with communicative acts. Other studies compared apes' reactive behavioural rate in the face of violations of requests expressed in manual points: Leavens et al. (2005) offered chimpanzees a high-quality and a low-quality food item choice, and varied the experimenter's response by having her deliver either the requested high-quality item (successful interaction), half of the high-quality item (partial success), or a low-quality item (failure). After the delivery, the chimpanzees persisted with gesture production more in the latter two conditions compared to when they received what they requested. Cartmill et al. (2007), similarly, found that orangutans showed more behaviours (indicating protest) when an experimenter shared a low-quality food item compared to when she shared a high-quality item (after their request for a high-quality item). In both cases, apes expressed their first-order intention towards a desirable food item, and reacted to the experimenter violating their request. It is not clear, however, whether the apes were reacting to not obtaining a desirable outcome, or whether they wer reacting to the experimenter's misaddressal of her specifically communicative intention. In other words, the Cartmill and Byrne (2007) and Leavens et al. (2005) studies were insensitive to the contrast found in Grosse et al. (2010). The protocols did not dissociate instrumental, external reference – first-order intention directed at a referent – and the communicative, internal reference – second-order intention directed at the recipient.

In this study, then, I sought to gauge chimpanzees' sensitivity to violations of their communicative requests. I dissociated communicative violations (caused by the experimenter) to either instrumental intentions (first-order, external), or specifically communicative intentions (second-order, internal). Specifically, I presented the chimpanzees with an object choice situation in which an array of high-quality (HQ) and low-quality (LQ) food items arranged openly in a line next to each other, available for choice. Similarly to other studies which tested the sensitivity to violations in interactions (Cartmill & Byrne, 2007; Leavens et al., 2005), I either met their choice with a successful interaction (sharing the item they chose), or with a violation (sharing a different item from the table) and compared their subsequent rate of behaviour between the conditions. Unlike the above reported studies, in the probe (violation) trials, I caused different types of violations.

In experiment 1, the array consisted of two adjacent pairs of HQ and LQ items. After the chimpanzees pointed at an HQ item, I shared either 1) the requested HQ item, which amounted to the 'baseline' successful interaction; 2) an adjacent LQ item, under the pretence of having mistaken the line of point, the 'honest mistake' condition'; 3) a distal identical HQ item, thus satisfying the apes' goal, but misaddressing her communicative intention, the 'identity mistake' condition; or 4) a distal LQ item, under the pretence of a deliberate offence to the chimpanzees's request, the 'offence' condition.

The crucial contrast of the design was between the baseline and the identity mistake condition, mirroring the logic of Grosse et al. (2010) whereby, in either condition, the chimpanzee got *a* HQ item, but in the honest mistake condition, it was not the one they requested, but another from the array equivalent in quality. In both cases, their external goal was met, but in the honest mistake condition, their communicative intention alone was violated. An alternative model would expect no difference in the behavioural rate between the conditions, and the inferential model would expect more behavioural rate in the honest mistake condition, similar to the effect in Grosse et al. (2010).

In line with the reviewed literature, I conjectured that the alternative model was more likely to explain chimpanzees' behaviour. However, its prediction is a non-difference, which is not ideal from a statistical position. To help with supporting the expected null effect, I included two conceptual replications of documented effects, both of which served as validation procedures for the experiment. First, the contrast in behavioural rates between the baseline success and the offence condition would match the effect in Cartmill and Byrne (2007) where the apes reacted more when they got the food they did not request, compared to one which they did request.

Second, I sought to replicate an effect akin to the unwilling-unable effect (Call et al., 2004). I compared the honest mistake and offence condition where, in both cases, the chimpanzees received the wrong LQ referent, but in the honest mistake, their actual point was adjacent to the chosen HQ item. The honest mistake could conceivably be construed as a misinterpreted point direction, but in the offence condition, I shared the item on the opposite side of the table. In the latter, mistaking the line of point was impossible, and therefore my sharing the wrong item could be construed as deliberate. In Call et al. (2004), chimpanzees protested more when, in closely matched actions, the experimenter acted being unwilling than when he acted being unable to share food items, suggesting sensitivity to the purposefulness of the experimenter's actions. Likewise, I expected a higher behavioural rate in the offence condition than in the honest mistake condition.

3.2. Experiment 1

3.2.1. Methods

Participants

Seven chimpanzees (*Pan troglodytes*) (six female) mean age = 37 (range 15 - 57 years) housed in Twycross Zoo with mixed rearing histories and limited previous experience with psychological tests. Recruitment was opportunistic and voluntary; I sat at the designated position of the bed area and waited for chimpanzees to approach. They were free to leave testing at any point and were not coerced into participation, with limited attempts to motivate them to approach by presenting food rewards.

Procedure and materials

The chimpanzees were tested in their bed areas with one door access to the indoor day area, and side access to adjacent bed areas, which means that chimpanzees from the group were freely roaming the testing area, and participation was voluntary and opportunistic. A mesh wall separated the bed area from the keeper area where I sat with their apparatus, against one of the mesh walls. The chimpanzees were able to leave the testing area (area against the testing table) without constraint from myself or the keepers. They were not food- or water-deprived and received their routine diet. The keepers would occasionally call or distract the chimpanzees to facilitate participation or discourage overcrowding (which could lead to distraction or fighting).

The overall testing setup involved a cooperative pointing and food sharing exchange between the chimpanzee and myself. The physical apparatus involved a 45cm tall folding table with approx. 70*40 surface and a sliding panel of approx. 50*40 surface.

Before entering the test phase, all of the participants underwent a preference test in which they chose between 2 food items, a 2*2*2 cm courgette, and a 2*2*2 cm sweet potato cube. Food type was chosen with input from keepers. In order to 'pass' a preference test, the chimpanzees had to show two subsequent sessions with at least 10/12 choices of the same type. The preference trials also served as training for object choice dynamics, seeing as the present group of chimpanzees had little-to-no testing experience.

The food pieces were spread in a 4-item array, grouped in two pairs of different items (one HQ sweet potato and one LQ courgette), one pair on each side. The pairs on each side were placed 1 mesh cell width (5cm) apart, whereas the closest items between the pairs were 4 cells width (20cm) apart. The food items were placed on the end of the sliding panel closer to the mesh. The sliding panel was pushed forward to signify the choice stage.

In a test session, I first placed a small occluder that partially obstructed the view of the food item array. After baiting the items in an array behind the occluder, I waited for the chimpanzee to retract any fingers from the mesh (protruding fingers more than three cells above the table were not considered as points and were allowed). I also used a predetermined gesture and verbal request to clear the mesh. Once the chimpanzees removed their fingers, I removed the occluder, pushed the sliding panel forward, and waited for the chimpanzee to choose a food item.

The array consisted of 4 items, or 2 pairs of adjacent different items, each positioned under a mesh cell. There were 4 conditions in total, and with 2 different array configurations (counterbalanced, each subsequent trial switched configurations). Food item configurations were organised thus (read iconically, left-to-right):

Configuration A: (HQ1)-(LQ1) — (HQ2)-(LQ2)

Configuration B: (LQ1)-(HQ1) — (LQ2)-(HQ2)

Numbers, here, designate position, but HQ1 was featurally indistinct to HQ2 (sweet potato pieces), as were LQ1 and LQ2 (courgette pieces). The food items within a pair were separated by one mesh cell (5cm), and the closest items between pairs were separated by four mesh cells (20cm).

Since the panel was open mesh, the pointing position was distributed relatively continuously. Thus, in order to read the chimpanzees' gestures, I used a predetermined set of rules to determine how to interpret the position of the point. The cell which was protruded by the most extended finger was considered the choice cell. Food items were placed on positions corresponding to mesh cells, so that a point through one cell would select its corresponding panel position in front of the cell. A protrusion through a cell between food items was ambiguous, and a protrusion through two central cells was uninformative. In these cases, I would delay food delivery and hold the panel in the forward position until the chimpanzee changed her point to an unambiguous position. In a minority of these cases, the chimpanzees angled their fingers substantially so that the tip of their fingers would be unambiguously directed towards a food item which would entail a choice. I would then deliver either the item pointed to, or another item, depending on the condition. Points that were positioned more than three cells above the plane of the table were not counted as choices, as they likely entailed leaning against the mesh.

The procedure assumes that the chimpanzee pointed to one of the HQ items, preference for which was established in the preference tests. In cases where the chimpanzee pointed to an LQ item, the test trial was reset. If the chimpanzee chose an LQ item in 3 subsequent trials, she was assumed not to be appropriately motivated or focused, and the session was stopped until the next testing day.

Once the chimpanzee provided a legitimate pointing gesture, the procedure forked according to either of the 4 conditions (illustrated in Figure 3.1):

1. Baseline success: I shared HQ1, the food item pointed at
2. Honest mistake: I shared LQ1, the item of the opposite quality but from the same pair as the pointed at item
3. Identity mistake: I shared HQ2, the item from the opposite pair and of the same quality as the pointed at item
4. Offence: I shared LQ2, the item from the opposite pair and quality

Figure 3.1. Apparatus and baiting configuration.

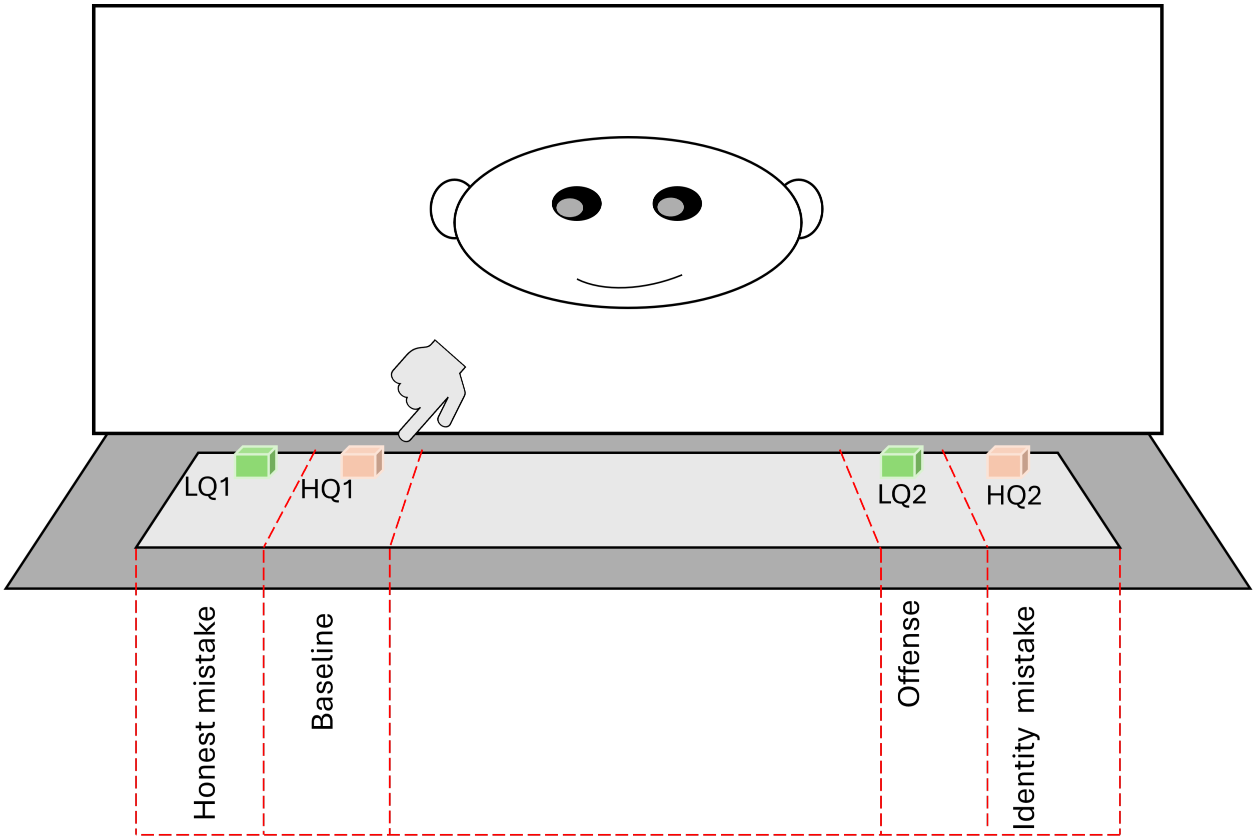

Note. The condition decided which food item I shared in response to their point towards an HQ item, and is shown in the figure below the items. Every trial assumed that the chimpanzees pointed to one of the HQ items, and determining which item I shared was contingent on which HQ item they pointed to.

After the choice, I reached for an item on the sliding table with tongs in my right hand, and with the left, pressed the button on a stopwatch at the moment when the tongs grasped the item. I delivered the item through the feeding tube, marked the coding sheet, and waited until 10-seconds elapsed. The remaining items were left on the table during this 10-second period. Then, I replaced the occluder and rearranged the food items for the next trial. Each item moved positions to one on the right, and the shared item was replaced. In invalid trials where the ape pointed at an LQ item, I reshuffled the food so that the HQ1 item switched position with HQ2, and LQ1 with LQ2. This way, the actual items were shuffled around, but the array stayed the same.

Trials which were interrupted so that they interfered with the crucial 10s period were considered invalid and were repeated. Examples include the participant being displaced, food reward being stolen, or background commotion causing the participant to become distracted. In cases where participation exposed the chimpanzee to undesirable attention, e.g., if a dominant continuously lurked around, waiting for an opportunity to steal food, the experimenter paused the session until the situation changed. Such interactions could otherwise lead to physical altercations.

A session consisted of 15 trials, 12 of which were baseline success trials. Either of the remaining 3 probe trial types appeared only once per session to keep the chimpanzees' expectations biased towards successful interactions. The trial distribution was arranged in a pseudo-random manner so that the probe trials were separated by at least 2 baseline trials, and the first 3 trials of any given session were necessarily baseline trials. Thus, every session started with standard successful interactions, and further in the session, an occasional 'violation'(probe trial) would happen.

Measurement

The recorded measurement were combined indices of frustration, protest, or repairing the communicative interaction by the chimpanzee during the 10 second period subsequent to food delivery. Due to a low overall rate of behaviour, the outcome variable was compiled into a dichotomous index of presence or absence of a behaviour in a given trial, therefore I compared the numbers of trials that had at least one behaviour between different trial types. Crucially, the behaviours had to be directed at the experimenter, food items, or apparatus. The index was conceptually based on the 'behavioural rate' from the unwilling-unable paradigm (Call, 2004) and gestural repair (Cartmill & Byrne, 2007), however, operationally, in this study, there were idiosyncrasies pertaining to chimpanzees' habits and testing situation and the fact that I tested them with the rest of the troop present, rather than separately. Most of the behaviours included were decided based on informal observation and interacting with the sampled chimpanzees. The ethogram with brief descriptions of behaviours is described in Table 3.2.

Table 3.2. Ethogram table of coded behaviour.

Coded behaviour	Description of the behaviour
Mesh touch	Touching the mesh with their hands, excluding cases where they held for balance or resting their hand
Point to the same side	A pointing action* (protruding hand through the mesh) that is on the same half of the mesh as the initial point
Point to the opposite side	A pointing action* that is on the side opposite to where the initial point was
Head nodding	Head nodding gesture; one bout marked as a single event
Clapping	Clapping hands together that is not clasping hands together
Manipulating the apparatus	Attempting to manipulate the table or sliding panel either by hand or stick
Mouth contact	Physically contacting the mesh with the mouth
Food grunt	Vocalisation usually occuring in anticipation of apetitive food

Note. *A pointing action was defined as a protrusion with the fingers not higher than three mesh cells (15cm) above the table. Position was determined as belonging to the cell that the most extended finger protruded.

Importantly, the ethogram was not set a priori in full. The initial ethogram included some behaviour that I later decided were not appropriate due to the group setting of testing: behaviours like scratching can be indices of frustration, but will necessarily correlate with orthogonal social factors like the proximity of a dominant conspecific. New behaviours like second pointing were introduced either during or after data collection for the simple reason that they were not anticipated before testing. Not all behaviours should be interpreted as indicating the same state, i.e., frustration, but all should generally be compatible with an inclination to repair or react to the interaction. Therefore, the combined index should relate to factors related to the interaction between the chimpanzee and experimenter during testing.

Behaviours were coded from testing videos offline using Solomon Coder 19.08.02.

3.2.2 Results

I used R version 4.4.2 (2024-10-31) for all analyses.

A separate coder coded 20% of the videos. Inter-rater agreement was moderate, $\kappa = 0.67$, $p < 0.05$.

Data is summarised descriptively in Table 3.3. Roughly half of the trials (49.09%) where the experimenter violated the chimpanzees' requests (sans baseline) had at least one behaviour recorded. This relatively low rate of behaviour should constrain any further inferences.

Table 3.3. Descriptive summary of the number of trials according to whether or not they contained a recorder behaviour, per condition type.

Trial names	**Sum of all behaviour events**	**N of trials with a behaviour**	**N of trials without a behaviour**	**% of trials with a behaviour**
Baseline	178	150	495	30.3%
Identity Mistake	16	15	41	36.59%
Honest Mistake	33	22	43	51.16%
Offence	36	25	42	59.52%

The main confirmatory analysis included three planned chi square contrasts with a Bonferroni correction. The family-wise type 1 error was set to .05, and the critical value for each chi-square contrast was 0.0167.

In the first 'validating' contrast, I compared the number of trials with at least one behaviour between the baseline success and offence conditions. There were more trials with a behaviour in the offence condition than the baseline condition, $\chi^2 = 4.37$, $p < 0.0167$, $\phi = 0.17$. This effect replicates the Cartmill and Byrne (2007) effect; chimpanzees reacted more to an unsuccessful and unpalatable outcome of an interaction than a successful palatable one.

The second contrast informed the 'validating' unwilling-unable-type effect replication; I compared the number of trials with a recorded behaviour between the offence and the honest mistake conditions. There was no significant difference between the conditions, $\chi^2 = 4.37$, $p = 0.04$, $\phi = 0.23$, failing to replicate evidence of an unwilling-unable-type effect.

The third, theoretically crucial contrast compared the baseline successful interaction with the identity mistake condition, which could support either the alternative model (no difference), or the inferential model (more behaviour in the identity mistake condition). There was no significant difference between the baseline and identity mistake conditions, $\chi^2 = 0.70$, $p = 0.40$, $\phi = 0.04$. Failing to find a difference alone is compatible with the alternative model (but the interpretation also hinges on the validating conditions).

Figure 3.2. Relative proportion of trials which had at least one behaviour recorded according to condition and those that had no recorded behaviour.

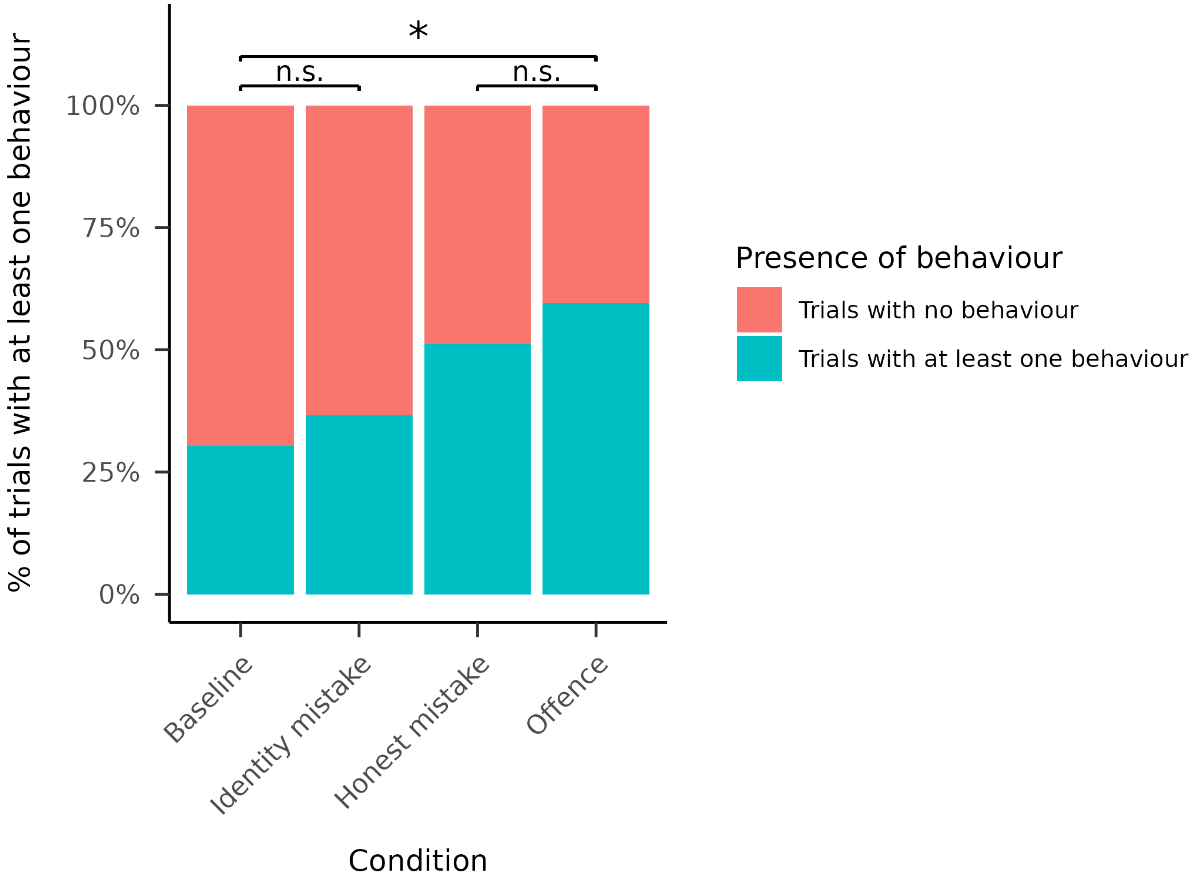

Note. The horizontal lines depict planned comparisons reported above. Significance level of all comparisons was set to 0.0167.

In a complementary exploratory analysis, I ran a mixed effects logistical regression with presence of behaviour as a binary outcome, trial type and trial order (across sessions) as fixed predictors, and subject as random predictor. I sought to test which conditions contributed to a difference in the number of trials with a recorded behaviour compared to the baseline condition. The family-wise error rate (0.05) was adjusted using the Bonferroni correction so that each comparison was weighed against $p = 0.0125$. Trial order (across sessions) did not contribute to the presence of behaviour ($p = 0.16$), suggesting no change across testing time. As for the conditions, the identity mistake condition did not yield more trials with a behaviour than the baseline intercept $\beta_{\text{imistake}} = 0.36 \pm 0.38$, z = 0.95, $p = 0.34$, while both the honest mistake and offence conditions had significantly more trials with a behaviour compared to the baseline (honest mistake: $\beta_{\text{hmistake}} = 1.07 \pm 0.37$, z = 2.89, $p < 0.013$; offence: $\beta_{\text{offence}} = 1.53 \pm 0.38$, z = 4, $p < 0.013$). Notably, the latter two conditions were ones where the chimpanzees received the LQ food item.

3.2.3. Discussion

In experiment 1, I found partial support for the alternative model of communicative intentionality which accommodates sensitivity to instrumental, but not communicative intention. I did not find a difference in recorded behaviours between when I purposefully delivered the not requested HQ item, as opposed to the requested HQ item (baseline vs. identity mistake). This failed in obtaining an effect similar to children correcting the misaddress of their request despite satisfying their goal (Grosse et al., 2010).

The protocol included an internal validation through replication of the effect of gestural repair in the face of request violations (Cartmill & Byrne, 2007) and the unwilling-unable effect (Call et al., 2004). The former obtained: there were more trials with a behaviour in the offence condition than in the baseline success condition. This effect is compatible with them being sensitive to an instrumental violation, but also to a generally unpalatable outcome (i.e., receiving an LQ item) regardless of the request.

The unwilling-unable-like effect (Call et al., 2004) did not obtain: I failed to detect a difference between the honest mistake and offence conditions. Possible reasons for the failure may be either group-specific factors like the rearing history of individuals, social dynamics, etc., or more tractably, idiosyncrasies of testing situations. In Call et al. (2004), the chimpanzees were in a passive role, only accepting routine food deliveries, while in the present study, they actively chose food items. In the former, too, they were tested individually, without interference from group mates, whereas in our study, the chimpanzees were tested in bed areas open to everyone from the group. Finally, it is possible that the present type of manipulation, which pertained to the position of referents, simply was not appropriate for stressing the purposefulness of the ‘offence’ (equivalent to unwillingness) of the experimenter.

Is partial validation enough to show that this testing setup can detect effects pertaining to communicative violations? It is possible that the unwilling-unable-like effect is not suited for this setup, and indeed, the original setup (Call et al., 2004) differs from the present one more so than the Cartmill and Byrne (2007) setup. It may be that the unwilling-unable effect is not a good validating procedure for the present testing situation.

An effect showing sensitivity to unmet communicative goals (Cartmill & Byrne, 2007) in general is not necessarily communicative in nature. Plausibly, the chimpanzees here reacted more simply to receiving *an* LQ item, not because it is a sign of a failed interaction, but because it is an unpalatable event. Findings from the exploratory model seem compatible with this notion: the conditions where I shared an LQ item were more likely to yield a trial with a behaviour compared to the baseline, but not also the condition where I shared the wrong HQ item. However, both of these conditions where I shared the LQ item can be construed as either instrumental violation or as a generally unpalatable outcome. A further control condition would be needed to eliminate the 'general palatability' sensitivity hypothesis. For example, one could introduce a condition where the animals would receive another type of LQ item that was not offered up for choice and would be administered by either another experimenter (not the recipient of the request), and in a non-social manner (e.g., by an automated mechanism). In such a condition, the event would not be contingent on their request (if the shared item was novel), and would not be socially mediated (if food were delivered inanimately). If such a condition provoked similar rates of reactive behaviour as an offence condition, this would support the idea that the chimpanzees are responding to the general palatability of an event, regardless of its social nature. As is, we cannot determine whether the chimpanzees' reactivity was due to sensitivity to an instrumental request violation (i.e., receiving the *wrong* and unpalatable item), or a general palatability sensitivity (i.e., to receiving an unpalatable item generally).

Importantly, the inter-rater reliability showed only moderate levels of reliability, $\kappa = 0.67$. Since both I and the independent coder coded the behaviours from the videos, it follows that the ethogram was not unambiguous enough for reliable discernment of behaviour. This should constrain any further inference.

The overall rates of trials with at least one behaviour were low: between the conditions with any type of violation, only 49.09% of the trials had a recorded behaviour. It is possible that the distribution of different trial types biased the chimpanzees away from being reactive. Within any given session, there were only 3 probe trials with a 'failed' interaction, and a session always opened with 3 baseline (successful) trials. The reasoning here was to prime the expectation of successful interactions, and hopefully exacerbate the surprise of the rare violations. The same arrangement could have likewise decreased the cost of a failed interaction – ending up with an occasional courgette is easy to endure when there is enough sweet potato for consolation.

In experiment 2, then, I attempted to inform roughly the same theoretical question, but I addressed the above concern regarding the distribution of probe trials. In the following design, I brought the ratio of probe and baseline trials within a session closer (1:2 instead of 1:5), and omitted the rule by which probes could not occur at the start of a session. I also changed the setup from a 4-item choice to a 2-item choice, hopefully making the choice simpler for the chimpanzees, which could plausibly make them 'commit' more to their choice, and by extension, protest more when a delivery failed. Both modifications intended to exacerbate the overall rate of reactive behaviour.

I placed two items to the left and right sides of the sliding panel. I varied the configuration that made up the choice, as well as how I responded to the chimpanzees' choices. The choice was either between two HQ items, or between an HQ and an LQ item (same/different identity factor). I either delivered the item that the chimpanzees pointed at, or the opposite item (regular/inverse delivery factor). Thus, the choice was either between two different or between two of the same items, and I either delivered an item which they requested, or the inverse.

A comparison of behavioural rates between the regular and inverse delivery conditions constituted different types of violation, and the type depended on the identity conditions. A contrast between the different item conditions would have shown the chimpanzees' reactions to not getting what they wanted, or because their request was misaddressed – in other words, *either* because the failure is instrumental or communicative. The second contrast between the same identity conditions informed the question of sensitivity to specifically communicative, and not instrumental failure: in the same conditions, the chimpanzees met their instrumental goal (obtaining an HQ sweet potato), but in the inverse delivery, their request was misaddressed (they received the opposite HQ sweet potato). The alternative model of intentional communication would expect no differences between the same conditions, but a greater behavioural reaction to a violation between the different conditions.

Similarly to experiment 1, the first contrast also acted as a validation procedure; it sought to replicate the Cartmill and Byrne (2007) effect where orangutans reacted with a greater behavioural rate to receiving the wrong LQ food item over the requested HQ one. Thus, the present testing situation can be validated by the replication, and so even a failure for a difference in the correct direction to reach significance is rendered more informative: one would be able to say that the crucial, second contrast effect size is at least not as large as that of the first contrast.

3.3 Experiment 2

3.3.1 Methods

Participants

The same 7 chimpanzees from experiment 1 participated in experiment 2, in the same voluntary and unrestricted setting and housing.

Procedure and materials

The same materials were used as in experiment 1, with the exception of the number of items available in a given trial. Here, the choice was between 2 food pieces: either between HQ and LQ ('different' conditions), or between two HQ items (‘same’ conditions), and I either delivered the item that the chimpanzee pointed to (‘regular’ conditions), or the opposite item (different ‘conditions’).

The design, then, was 2*2: food identity (same/different) * food delivery (regular/inverse):

1. Same regular: choice between two sweet potatoes, I shared the correct item
2. Same inverse: choice between two sweet potatoes, I shared the inverse item
3. Different regular: choice between a sweet potato and a courgette, I shared the correct item
4. Different inverse: choice between a sweet potato and a courgette, I shared the wrong item.

Each session consisted of 12 trials: 4 same regular; 4 different regular; 2 same inverse; 2 different inverse. The probe ‘inverse’ conditions occurred more frequently than in experiment 1, but were still in a minority. All trials were distributed evenly and pseudo-randomly across the session, i.e., the same trial could not repeat more than 2 times subsequently. I also pseudo-randomised the order of baiting, whereby no more than 3 subsequent trials were baited at the same side first.

The rest of the procedure pertaining to the order of events and coding was the same as in experiment 1:

I first set the small cardboard barrier on the sliding panel, baited the two food items using plastic tongs (left-right order was counterbalanced), and slid the panel forward for the chimpanzee to choose between. After receiving the choice, I pulled the panel back, grabbed an item, marked the beginning of the critical 10s period, shared the food, marked the trial, and waited for the 10-second period to elapse, which ended the trial. The other, not chosen item was left on the panel during the 10-second period. I then replaced the cardboard barrier and prepared the subsequent trial.

I coded the same behaviours as in Experiment 1 (see Table 3.2) using Solomon Coder 19.08.02 and analysed the presence or absence of a behaviour in a given trial between conditions.

3.3.2 Results

I used R version 4.4.2 (2024-10-31) for all analyses.

A separate coder coded 20% of the videos of experiment 2. Inter-rater agreement was moderate, $\kappa = 0.59$, $p < 0.05$.

Across all trials, I recorded at least one behaviour in, on average, 24.28% of trials, and 34.97% of trials across probe ‘inverse’ trials. Thus, I observed a behaviour in less trials than in experiment 1 (49.09%), suggesting that the modifications to the design did not provoke more behaviour, as intended. This low rate of observed behaviour should constrain further inferences, the chimpanzees may not have been inclined to correct their behaviour.

Table 3.4. Descriptive summary of trials with respect to whether they had a recorded behaviour, per different trial types.

Trial Names	**Sum of all behaviour events**	**N of trials with a behaviour**	**N of trials without a behaviour**	**Proportion of at least one behaviour**	**Proportion of at least one behaviour per delivery type**
Same Regular	30	24	79	23.3%	
Different Regular	4	4	99	3.88%	13.59%
Same Inverse	20	19	32	37.25%	
Different Inverse	29	17	35	32.69%	34.97%

The family-wise type 1 error was set to 0.05, and Bonferroni correction set the critical p-values for both planned contrasts to 0.025. First, comparing the 'different items' trials (regular and inverse) revealed more trials with a behaviour in the inverse than the regular conditions, $\chi^2 = 24.49$ $p < 0.025$. When the chimpanzees requested the HQ item, they showed more behaviour after the I shared the LQ item, compared to when I shared the HQ item. This effectively replicates the Cartmill and Byrne (2007) effect.

Second, the theoretically crucial comparison between the 'same' trials failed to reveal a significant difference between the inverse and regular conditions, $\chi^2 = 3.30$, p = 0.07, which fails to suggest that the chimpanzees were sensitive to the specifically communcative violation. The relative frequencies are graphically summarised in Figure 3.3.

Figure 3.3. Relative proportions of trials that had a recorded behaviour per condition.

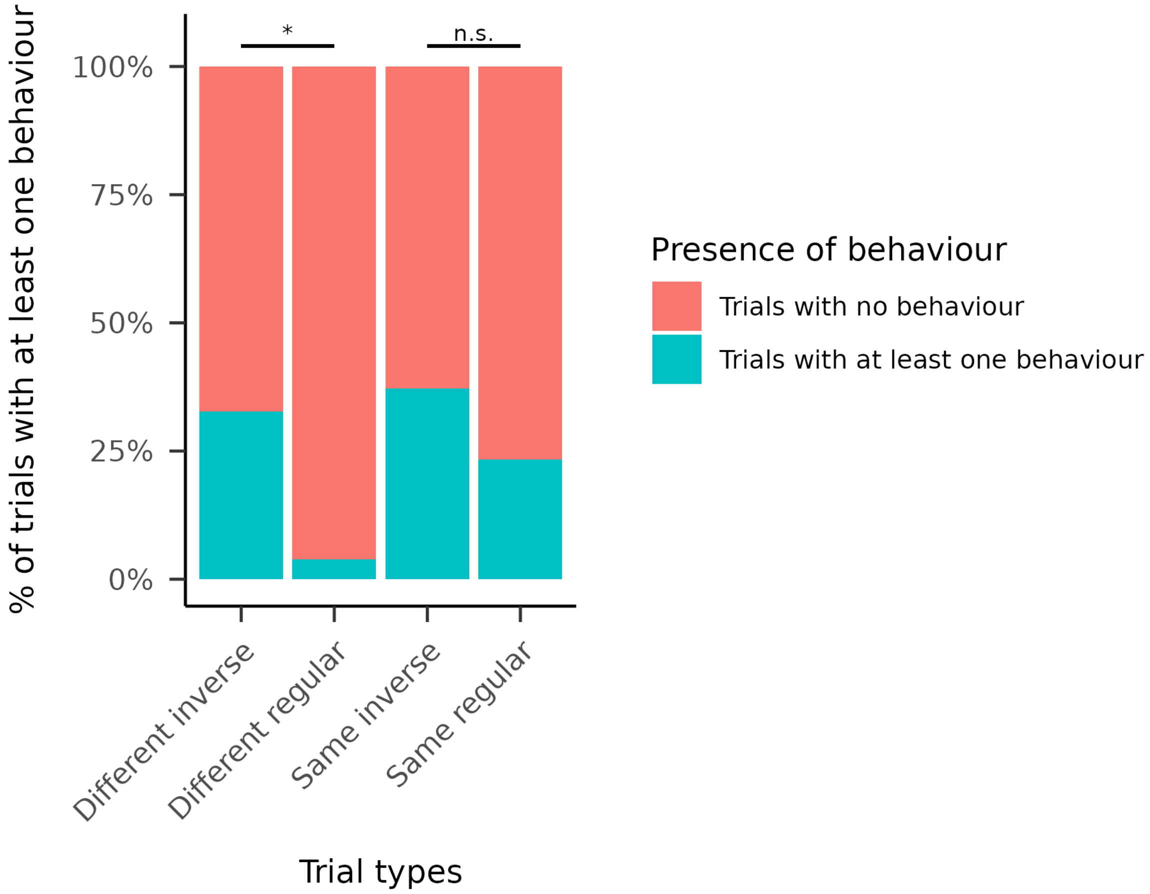

Note. Horizontal black lines designate planned comparisons reported above. Critical type 1 error set to 0.025.

3.3.3 Discussion

The proportion of critical trials with at least one behaviour recorded in the face of communicative failure was lower still than in experiment 1 (49.09 % vs. 34.97%). The new design implementations did not increase reactions to violations of requests. Between two contrasts, I found data compatible with the idea that chimpanzees protest to instrumental failure in the interaction (getting the wrong item of poorer quality), again replicating the Cartmill and Byrne (2007) effect, but not also in specifically communicative failure (getting the wrong item of an equivalent quality). Like in experiment 1, however, it is also possible that the chimpanzees were reacting to receiving an unpalatable item regardless of the social nature of the interaction.

It is not clear what aspect of the design contributed to the lower behavioural rate across trials in experiment 2 over experiment 1. It is possible that by introducing more 'irregular' trials where I delivered the inverse food item provoked a change in the chimpanzees' expectations about my cooperativeness, which, in turn, made them less sensitive to the violations. This would run counter to the intention of increasing the cost of, and therefore vigilance to, an unsuccessful interaction by increasing their relative frequency.

3.4 General Discussion

I set out to examine chimpanzees' sensitivity to communicative violations pertaining to their specifically communicative intentions, as opposed to instrumental intentions. Across two experiments, I found no support for the notion that the chimpanzees were sensitive to specifically communicative intention violations, which I take to indicate that they are unlikely to produce gestures with second-order intentions, i.e., intentions to communicate directed at the recipient (me) rather than at the overt target in the world (food items).

In contrast, I found some data compatible with the idea that the present sample was sensitive to instrumental violations, i.e., when the chimpanzees received the wrong and unpalatable food item: across both experiments, I found more trials with a behaviour subsequent when the chimpanzees got the wrong LQ item than the correct HQ item. This is compatible with the Cartmill and Byrne (2007). In experiment 1, further, the logistical regression model revealed that receiving any (wrong) LQ item made it more likely to observe a trial with a behaviour. There was no such increase detected in when the chimpanzees received the wrong HQ item.

The present data is unable to disambiguate between two hypotheses about what drove the chimpanzees' behaviours in the face of request violations: either they were responding to receiving an unpalatable item, or they were responding to their my failure to meet their request. The former alternative need not be communicative or social in nature. In future studies, it would be informative to include a control condition where the requesting animal receives an LQ item without the recipient of the request delivering it herself (so that it is shared by another experimenter, or by some automatic means). Should this control show the same behavioural rate as one where the request is met with delivering an LQ item – offence or different inverse, in experiments 1 and 2, respectively – then their behaviour is likely driven by the unpalatability of the event, rather than a communicative violation.

A sensitivity to instrumental, but not specifically communicative violations is in line with what I referred to as the alternative model of intentional communication, by which intention holds primarily environmental ('external'), and not mental ('internal') information. This model builds on the interpretation in Kano et al. (2018) who found that great apes are sensitive to some degree of ostensive communicative intention, but only in so far as it orients them to seek information in the external world. From the recipient's side, we may expect the communicative intention in the alternative model to act as a trigger for searching the external space for plausible referents. In the inferential model, however, recognising the communicative intention would act as a trigger for inferring the meaning held in the producer's mind, and only secondarily search for a candidate in the environment. To make this distinction more operationally tractable, consider Anne pointing to cloudy sky and uttering 'It looks like a dog!' for Sally, who will search the sky for a likely target of her utterance. Searching the space for candidate referents alone is compatible with either the inferential or alternative models. However, the search *strategy* might dissociate the two models: in the inferential model, Sally is primed by Anne's representation of a dog-shaped object, which effectively pre-constrains her search for a candidate in the environment. She might stop at the first sufficiently dog-like cloud she sees. Following the alternative model instead, though, she might seek out all dog-like clouds, and only after that, select one. The mark of the inferential model is that the recipient first attends to the internally determined meaning in the signaler's mind, and only then searches the environment. The alternative model starts by searching the environment, and only then matches it with the signaler's intent. This thought experiment is empirically tractable in a further design: setting up an interactive situation so that a signaler gazes or points ambiguously at a group of different objects - some of which look similar - would the recipient assume that she is singling out one object, or any object of equivalent features? Plausibly, the inferential model would pick out one referent as the likely target, and the alternative, all similar objects.

Returning to the design at hand, we may offer some tentative support for the alternative model. For a more complete design, one would need to include both a) a control for sensitivity to general palatability of an outcome, and include a population which presumably possesses a sensitivity to specifically communicative intention, e.g., older children. Presumably, an 'able' population would show sensitivity to specifically communicative intention in a similar design, unlike the chimpanzees presently, which would deliminate the 'upper limit' of intentionality. In conjunction with a 'lower end' control of alternative LQ administration (either by a new experimenter or by physical automation), the protocol would be opportune for informing intentional communication.

Chapter 4: Communicative disambiguation

4.1 Introduction

An animal is acting intentionally, or holds a mental state about something, when she refers to something outside itself (Gómez, 2009). An intention or a mental state can manifest behaviourally as a manual point towards an object O, or a verbal utterance about O, and in both cases, a gesture is *about* O, i.e., O is the content or referent of her mental state. In communicative situations, the intentional act is enmeshed within a more complex situation where an organism is not only directed at the object but also at another agent. Under some formulations of intentional communication, signallers use communicative acts in order to convey their intentions *for the benefit of* others or with the goal of modifying a recipient's behaviour in a way conducive to their goal (Wilson & Sperber, 2012). A manual point directed at a glass of water for the benefit of a recipient is likely an indication that the producer wants the glass, but it is also directed at the recipient, in trying to make her aware of this. We know that, for a successful communicative interaction, it is not sufficient to just refer to the object implied in our goal (e.g., a glass) – if this were true, any act of attending to some object in the world would qualify as communicative. Intentional communication can be independent of such an object: even if a signaller fails to find the right words, or locate the item in space accurately, the recipient may still find her attempt meaningful. If my friend is explaining a complex idea and struggles for an appropriate word, I can still appreciate her intention of explaining something to me, without knowing exactly what she means; recognising her intention, I might help her out by inferring what she really means. Thus, recognising the intention to communicate alone, even without identifying the correct referent, is sufficient for a successful communicative interaction. So goes, roughly, Wilson and Sperber's (2012) reading of Gricean intentional communication which highlights the psychological importance of producing intentional communicative acts, and recognising this intent in others. This type of intention need not have a straightforward behavioural marker, neither from the recipient's nor the producer's side, especially if a straightforward reference to an object is out of the picture. In this reading, an intention to inform the interlocutor is necessary, and possibly sufficient for a meaningful communicative interaction.

A Gricean reading like this involves a complex coordination of mental states between interlocutors, which led some comparative psychologists to consider alternative, leaner formulations of intentional communication. Townsend et al. (2017) provided three conceptually and operationally tractable criteria for intentional communication to serve as alternatives to a Gricean reading. For an act to count as communicatively intentional, 1) a signalling animal S should have a goal pertaining to content O; 2) S should produce the goal voluntary and recipient-directed signals conveying O; and 3) that signal should cause change in the recipient that is relevant to O. These conceptual markers constitute an alternative to readings like Wilson and Sperber's (2012) which seem to require an additional component of specifically communicative intention 'on top' of an intention towards something in the world. The Townsend et al. (2017) model is workable for interactions specific to nonhuman primates. For example, the three markers map well onto typical interactions of great apes in a captive experimental setting: a chimpanzee may desire a food item on the table (1), produce a manual point for the experimenter (2), who will then share the food and satisfy the signaller chimpanzee's goal (3). This interaction is compatible with Townsend et al.'s (2017) criteria, but not unambiguously and conclusively so. To better inform the model, one may show that apes are specifically sensitive to these indices. An animal must show that her behaviour co-varies with 1) changes in the some content embedded in the signaller's goal, (2) availability of a recipient's attending to the signal, and (3) the behaviour the signaller induced in the recipient. I will weigh each of the criteria against available evidence of great apes' communicative sensitivities.

First, I will consider the second condition pertaining to recipient-directedness. If an animal is concerned with conveying something to someone, she must consider factors pertaining to their readiness to receive the message, such as bodily cues of attentional states. Several studies examined the apes' proclivities to adjust their manual and vocal gestural behaviour depending on the attentional state of the (usually human) recipient. Hostetter et al. (2001) showed that chimpanzees were more likely to produce vocalisations when the experimenter was turned away from them, but more manual (visible) gestures when she was oriented toward them. In Hostetter et al. (2006), they produced more visible gestures towards a human whose eyes were open than towards one whose eyes were closed, as well as when the human's mouth was covered than when her eyes were covered. Call and Tomasello (1994) showed that two orangutans (one enculturated, one captive) provided more communicative gestures when an experimenter was attentive, compared to when he was not. The former also showed sensitivity to the experimenter's knowledge about food or tool location. In an observational study with conspecific interlocutors, chimpanzees produced no visual-only gestures (e.g., arm-raise) when the recipient was not looking at the subject, while auditory-visual (e.g., ground-slap) or tactile (e.g., poke) were more frequent when the recipient was not looking at the subject (Tomasello et al., 1994). Povinelli et al. (1996) showed that chimpanzees gesture towards a food item that an experimenter is looking at or manipulating, even when the item is an undesirable food piece, and the desirable one is adjacent, which suggests that they prioritised the recipient's availability over the object of their goal itself. Poss et al. (2006) reported that orangutans and gorillas produce more manual gestures for an experimenter who is facing towards them compared to one who is facing away. They were also more likely to leave the testing area for the same reason. Thus, there is evidence of apes tracking cues of attentional and perceptual states of communicative recipients, and of varying their gestural production accordingly.

Not all body orientation cues are treated the same, however. Kaminski et al. (2004) showed that the rate of apes' communicative behaviour will increase when the experimenter's body is facing the apes compared to when she is either absent or facing away, although it did not seem to matter to them whether or not she had her eyes open or closed. The apes also seemed sensitive to face orientation when the experimenter's body was turned towards them, but not when her body faced away from them. This interaction effect suggests that the apes may have differentiated between perceptual and bodily affordances: when the experimenter is sitting with her back turned, she cannot share any food, so it may not matter to the apes whether or not she is looking, because she cannot share food anyhow. When she is facing them, however, it makes sense to gesture more when she can see, thus increasing the chance of a successful interaction. Bulloch et al. (2008) reported a similar dissociation between body and face direction sensitivity. This appreciation of bodily vs. perceptual affordances seems compatible with apes reacting differently to failures in food sharing interactions in which the experimenter delays food delivery by acting unwilling, as opposed to when he is unable to share food (Call et al., 2004); in both cases, the contrast is between their partner's ability and disposition to satisfy a communicative goal. Thus, apes show a degree of appreciation for both the attentional and perceptual affordances of their communicative partners, plausibly satisfying the second Townsend et al. (2017) criteria of recipient-directedness.

The third of Townsend et al.'s (2017) criteria of intentionality entail sensitivity to the behavioural change that signals cause in the recipient's behaviour. Operationally, this is relevant to cases of success or failure of communicative interactions – if I ask my friend for a glass of water, I may be disappointed if she pours me tea, to which I may communicate my dissatisfaction, or reiterate my request. Apes showed a degree of sensitivity to violations of their requests: Leavens et al. (2005) offered chimpanzees a choice between a banana, which was a high-quality reward, and a small pile of chow, a comparatively low-level reward. After the chimpanzees made their choice (they had a distinct preference for bananas), the experimenter delivered either a full banana (successful interaction), half a banana (partial success), or pieces of chow (failure). In the post-delivery phase, the chimpanzees persisted in gesturing more in the partial success and failure conditions, compared to the success conditions. Seemingly, they persisted when the outcome of their signalling was inadequate, i.e., not conducive to the goal of their supposed communicative intent. Cartmill et al. (2007) found a similar pattern of reaction to an unsatisfactory food-sharing interaction: when orangutans chose a desirable food item, but received an undesirable one – i.e., the interaction failed – they produced more gestures than when they received an undesirable one, i.e., when they reacted to a failed interaction.

These examples seem to meet Townsend et al. (2017) criterion of sensitivity to the behavioural outcome caused by their gestures. Their requests are presumably produced with the aim of getting the recipient to do something congruent with a goal.

I chose to discuss the first of the above-mentioned criteria because this is the target of the present study. Under this criterion, a producer must hold a goal pertaining to some content. This is the aspect of a communicative act which determines the content of an utterance, a point, etc. In cases where one is intending to something external, we may speak of straightforward Brentanian intentionality, by which an organism refers to something outside herself (Gómez, 2009), and the relationship between the agent and her object is a function of aboutness, i.e., the point is about the object. An ape may point toward a food item; therefore she intends to it, and the point is about the item. This applies to both the signaller and the recipient – both will try to identify the correct referent of a communicative act.

From the recipient's side, we can look at apes following point and gaze of another agent. In such situations, a subject has a chance to track the direction and point of termination of an agent's imaginary line of gaze or point and inspect what they are looking at. Psychologically, the situation is ambiguous – is the subject orienting in line with the general direction of an imaginary line, or is she expecting the agent to intend towards something? The latter situation can be said to be intentional, but a simple co-orientation is not. Studying gaze following, Povinelli and Eddy (1996) made this distinction in early gaze following experiments where they showed that chimpanzees follow a human's gaze not only in the direction of the line of sight, but also consider the physicality of opaque barriers, so that they do not expect the target of gaze to be behind an opaque barrier, from the recipient's perspective. When the human gazed towards an opaque barrier, the chimpanzees followed the gaze to where it met the opaque barrier, rather than where the line of gaze would terminate past it (Okamoto-Barth et al., 2007; Povinelli & Eddy, 1996; Tomasello et al., 1999, 2001).

A further clue about whether following gaze entails an intentional act can be gauged from the fact that apes extinguish their gaze following response across a few trials of identical demonstration (Call et al., 1998; Povinelli & Eddy, 1996; Tomasello et al., 2001), and they provide double looks, i.e., they look back at the experimenter after the first following response (Bräuer et al., 2005; Call et al., 1998; Okamoto-Barth et al., 2007). Both the extinction rates and the double looks suggest that apes may approach the situation with an expectation of the gaze to terminate at something informative, and when it is shown not to be (either by repeatedly gazing at the old target or missing a target entirely), the response is extinguished. Moreover, Tomasello et al. (2001) showed that, while both infant and adult chimpanzees followed an experimenter gaze, only adults extinguished their gaze following response across trials. Brauer et al. (2005) also reported that infants were less likely to show double looks compared to adults.

Further, MacLean and Hare (2012) exploited the idea that excited emotional responses to stimuli can convey novelty of information: across three experiments, the authors varied an experimenter's exposure to an object (therefore varying whether it was novel to them or not) and subsequently deployed three versions of the gaze following paradigm. In experiment 1, the actor emoted and looked at either a familiar or unfamiliar object, and the apes' responded by following the gaze either to the target object or past it, following the line of sight more distantly. The chimpanzees, but not bonobos, looked past the object more often when it was familiar to the experimenter compared to when it was novel to her, compatible with the idea that they expected something other than a familiar item to be the target of a surprise reaction. In experiment 2, both species searched behind an occluded space for a second object more frequently when the actor emoted towards the objects and was familiar with them, compared to when they were new to her. Thus, great apes show a proclivity to perceive other agents' gaze as conveying intentional attitudes towards the world: they expect it to terminate at a relevant target, which need not be immediately present perceptually, and this type of intention understanding may develop ontogenetically. This is compatible with the notion that great apes appreciate that other agents intend towards things in the external world, i.e., that their acts are intentional in a Brentanian sense (Gómez, 2009).

On the production side of things, one might look at whether and how communicative behaviour co-varies along levels of ambiguity in a communicative context. If the satisfaction of a (communicative) goal is contingent on the interlocutor responding appropriately to a communicative act, then one must ensure they get the message right; in noisy live interactions, the content embedded in the intentional act (e.g., a point) must be guarded against the risk of misinterpretation, or else the interlocutor cannot fulfill the signaller's goal. In a busy room at a party, I might greet a friend across the room and manually gesture to them to meet me at the door, and I might over-exaggerate my gestures in order to disambiguate my intent so to convey something like 'not the bathroom door, the bedroom door further down'. By over-exaggerating, I put additional effort into being specific so that it should be clear that I intended to communicate my target referent (bedroom door), and not another plausible candidate of my reference (bathroom door). Thus, by referring to specifically the target, and not the distractor − by disambiguating the referent − I provided maintenance of the gesture's meaning.

For an animal to amount to this type of specific disambiguation, at the very least, one should show that communicative production is sensitive to the features or spatial position of candidate referents. Gonseth et al. (2017) showed that chimpanzees' pointing changes based on the spatial distance of a food item: an item was placed either on a proximal or distal position away from the chimpanzee, and she requested it by pointing through a mesh wall. Their gestures were placed in a higher position on the mesh more often when the food was farther away than closer, which is compatible with them tracking the change in the item's location. The type of gestures also varied between conditions where the experimenter was present and, in the control, where she was not: the former yielded more visual gestures like manual pointing, whereas the latter yielded more audiovisual gestures like banging, clapping, or vocalisations. These results show that communicative acts vary along with the spatial position of the referent. This alone, however, does not exhaustively describe proactive disambiguation of specific content; in order to show that, one must construe a situation in which the signaller disambiguates the correct referent *against* a plausible, but incorrect referent.

In line with this demand for specific disambiguation, Tauzin et al. (2020) tested whether great apes are sensitive to a risk of a human recipient misconstruing their manual point for an incorrect referent. In an object choice situation where an experimenter shared a food item that an ape pointed towards, a high-quality (HQ) and a low-quality (LQ) food item were placed perpendicular to the axis of the panel separating the experimenter and ape. In experiment 2, the position of the food was placed either so that the HQ item was in front, closer to the mesh (HQ front condition), or so that the LQ item was in front (LQ front condition). If the signaller ape would point centrally, in the middle of the panel, the imaginary line of point would intersect both food items. In the HQ front condition, she would intersect the HQ referent first, but in the LQ front condition, it would be the LQ item. Pragmatically, the recipient is licensed to interpret the point to refer to the item that the pointing line intersects first and share that item. In the LQ front condition, then, the signaller would have to modify their point by positioning it away from intersecting the LQ item – e.g., pointing from the side or above, 'laterally'. Apes provided more such lateral points in the LQ front condition than the HQ front condition, which is compatible with the notion that they were disambiguating the referent for the recipient. I will refer to this effect as the laterality effect: more lateral pointing in the LQ front condition compared to the HQ front condition.

Other results from the same study run counter to the disambiguation interpretation: first, a control condition where only the HQ item was placed in the back position – the No food front condition – also yielded more lateral points than the HQ front condition. This means that the laterality-like effect was obtained without the risk of misinterpretation. Second, in experiment 3 of the same study, the authors expanded the test by varying the experimenter's knowledge state. The food items were now occluded by upturned cups, and the experimenter that had shared the food was either the baiter (knowledgeable condition), or an assistant (ignorant condition). The rate of lateral points did not differ between conditions, indicating a lack of sensitivity to the experimenter's mental state. Finally, the rate of lateral points in the LQ front condition was only 35.71%, which may be interpreted against the idea of a proactive intervention. Presumably, a situation that introduces ambiguity as a factor can be regarded categorically; a low rate of lateral pointing might not support that notion.

Given those considerations, it is possible that apes were not proactively disambiguating between referents for other agents, but only tracking and matching the spatial position of referents, or one referent. However, it is also possible that the methodological setup was not sensitive enough for the apes' disambiguation to surface. In the present study, I aimed to exacerbate the communicative ambiguity of the interaction by further varying the spatial arrangement of the food items. Building on the logic of teasing out the laterality effect between the LQ front and HQ front conditions, I introduced another spatial arrangement in which I placed the items closer together, almost adjacently. This should exacerbate the communicative ambiguity of the situation because the imaginary line of point is at greater risk of intersecting a distractor when it is closer to the target. Therefore, experiment 1 was a 2*2 design where I varied the relative position of the referents (HQ front/LQ front) and the proximity of the candidate referents (far/close). The far positions were equivalent to the setup in experiment 2 of Tauzin et al. (2020), while the close position conditions were a novel addition. Like the original study, I measured the proportion of modified lateral points between LQ front and HQ front conditions to gauge the laterality effect.

Should the laterality effect be driven by communicative disambiguation, it would be greater in the close conditions compared to the far conditions because the former are at more risk of misinterpretation. Introducing the new conditions also offers the opportunity to test a competing hypothesis, which proposes that the laterality effect is caused not by disambiguation, but by tracking the spatial position of a target referent. It is possible that the target referent, HQ item, exerts a 'gravitational pull' on the position of the points, whereby the closer it is positioned to the mesh, the more central the apes' points are positioned, and inversely, the farther it is positioned, the more lateral the position. This 'gravity' mechanism differs from communicative disambiguation because it is not communicative in the narrower sense (it is not for the benefit of the recipient), but a reaction to the position of a salient, HQ object. Effectively, the gravity hypothesis describes a distance estimate.

While the disambiguation mechanism, and hypothesis, sees the proximity factor as a measure of ambiguity risk, the gravity mechanism observes only the distance between the HQ item and the mesh. A disambiguation hypothesis would expect that an increase in proximity between the items (in the close conditions) should lead to a greater laterality effect, proportional to the risk of misinterpretation. The gravity hypothesis, inversely, would predict that the laterality effect should be smaller in the close conditions than in the far conditions since the distance from the HQ item and the mesh differs less in the former than the latter. More succinctly: the disambiguation hypothesis expects a greater laterality effect in the close conditions than the far conditions while the gravity hypothesis expects a greater laterality effect in the far conditions than the close conditions. Figure 4.1 illustrates maps these predictions graphically.

Figure 4.1. Graphical summary of the predictions of the gravity and disambiguation hypotheses.

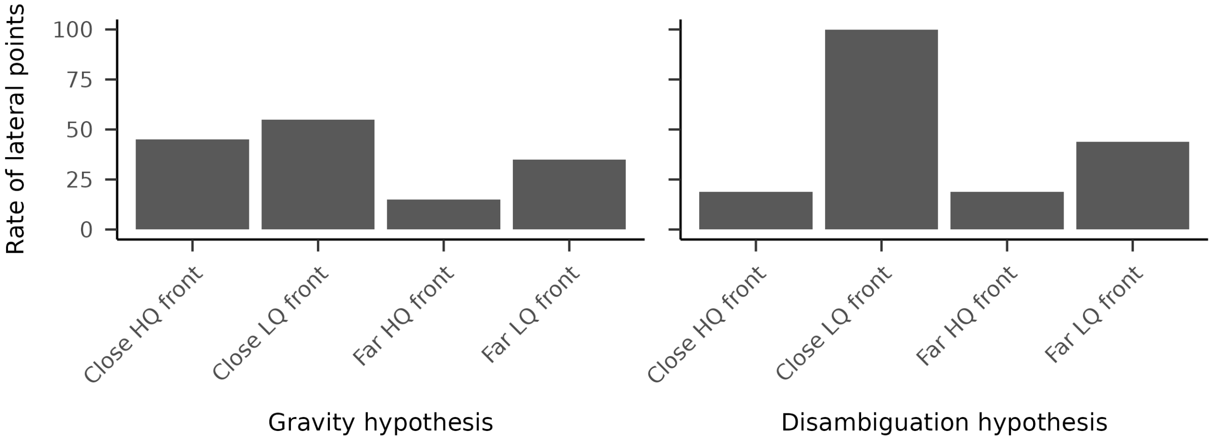

Note. The y-axis presents rough graphical estimations of relative positions for illustrative purposes only. Estimates are not of an exact numerical range.

4.2 Experiment 1

4.2.1 Methods

Participants

I tested the same sample as reported in Chapter 3: Seven chimpanzees (*Pan troglodytes*) (six female) mean age = 37 (range 15 - 57 years) housed in Twycross Zoo with mixed rearing histories and limited previous experience with psychological tests. Recruitment was opportunistic and voluntary; I sat at the designated position of the bed area and waited for chimpanzees to approach. They could also leave testing at any point and were not coerced into participation, with limited attempts to motivate them to approach by presenting food rewards.

Materials and Procedure

The housing conditions were the same as described in Chapter 3: the chimpanzees were tested in their bedding area, with free access to other bed 'rooms', as well as the indoor and outdoor areas. Participation was voluntary and there was no food or water deprivation.

I used a folding table and a sliding panel to present and offer the food items: approx. 45*70*40 cm table positioned against the mesh separating the keeper chimpanzees' areas. The sliding panel on which the food was placed sat on top of the table, approx. 50*40 cm. A small cardboard barrier was placed on top of the panel ahead of baiting to punctuate the trial; it did not occlude the food items but only served to mark the beginning of the trial and discourage premature food requests.

The quality of items was informed by keepers' diet observations and verified in preference tests, in which the chimpanzees chose between the HQ and the LQ food items – a sweet potato and a courgette piece, respectively – baited on the left and right sides of the sliding panel. A chimpanzee 'passed' if she showed a preference (for a sweet potato), i.e., 10/12 trials of the same choice, in two subsequent sessions. A chimpanzee continued to undergo preference tests until reaching criterion, or until it was no longer practical to test her. All chimpanzees that were included in the test stages demonstrated their preference for the chosen HQ food item, sweet potato, over the LQ item, a courgette.

The food items used in the preference and test sessions were 2*2*2cm cubes of courgettes (LQ items) and sweet potatoes (HQ items). I cut the pieces to size each morning of the testing sessions.

In the test phase, each trial started with me placing a small cardboard barrier on the sliding panel and encouraging the chimpanzees to remove their fingers from the mesh, if any (using a predetermined manual and verbal gesture). Once the mesh was clear, I baited the sliding panel. Two items, one HQ and one LQ, were baited in a line orthogonal to the mesh, either in a 'close' or 'far' positions (separated either approx 35 cm or 2 cm, respectively), and either so that one item was 'in front' closer to the chimpanzee and the other farther, or inverse. This yielded 4 conditions, illustrated in Figure 4.1:

1. Far HQ Front: items at the far ends of the panel, HQ closer to the chimpanzee
2. Far LQ Front: items at the far ends of the panel, LQ closer to the chimpanzee
3. Close HQ Front: items at the centre of the panel, HQ closer to the chimpanzee
4. Close LQ Front: items at the centre of the panel, LQ closer to the chimpanzee

Figure 4.2. Graphic illustration of the experimental setup for each of the four conditions of experiment 1

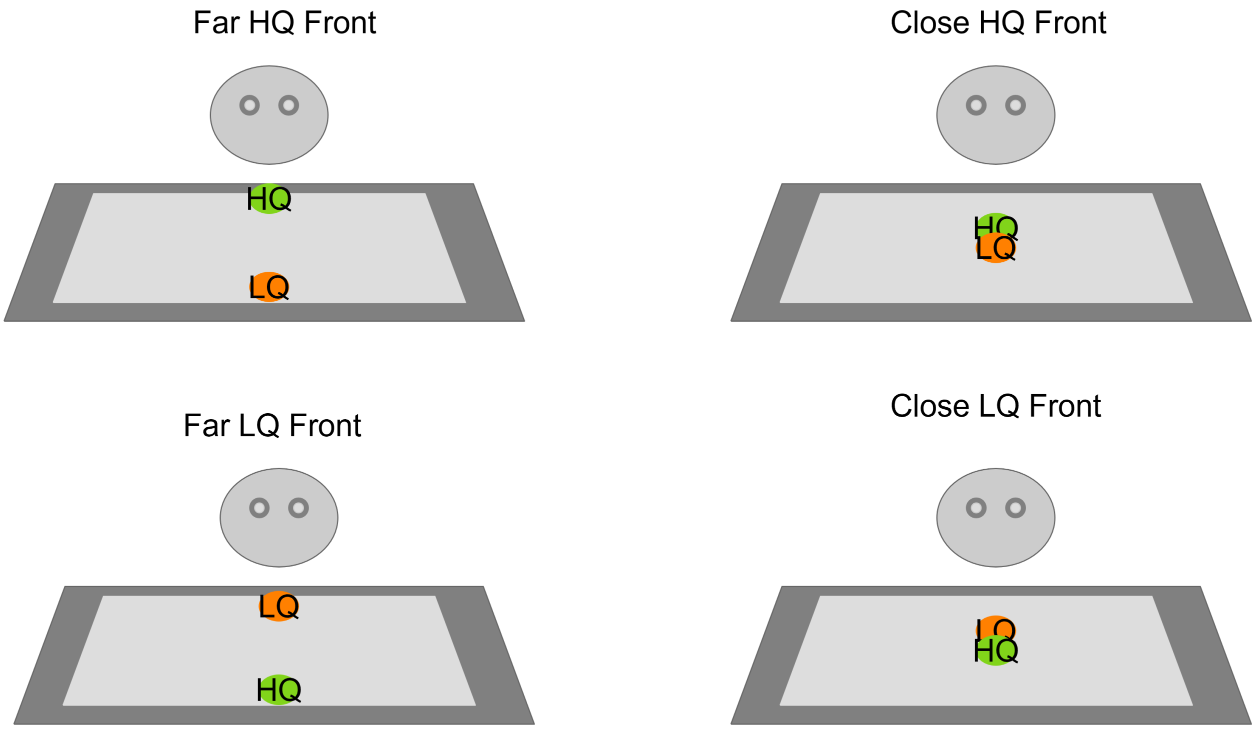

Far HQ Front

Close HQ Front

Far LQ Front

Close LQ Front

Each chimpanzee went through 4 sessions of 12 trials, 3 for each trial type. Trial types were pseudo-randomly distributed so that no more than two consecutive trials of the same type were permitted.

The location of the chimpanzee's point was the outcome variable of the test. Based on Tauzin et al. (2020), I categorised the points either as central, or lateral points. A central point was taken to be one positioned in the central 2*3 (W*H) cells of the mesh, and the lateral points were those outside of this area. I responded to a central point by sharing the item closer to the mesh, and I shared the item farther from the mesh as a response to a lateral point.

Since the chimpanzees often pointed with an open hand, protruding the mesh with several fingers at a time. I considered the mesh with the most extended finger as indicating the true position of the point. The central 6 mesh cells (central area) were marked with a permanent marker from my side only to help categorise the points' locations live.

At the end of each session, I rewarded the chimpanzee with an additional HQ food item, performed the routine gesture that signals the end of testing for that day, and refused to engage further.

4.2.2 Results

I used R version 4.4.2 (2024-10-31) for all analyses.

A separate coder coded 20% of the test sessions, arriving at high inter-rated reliability, $\kappa = 0.83$, $p < 0.05$.

Compared to the original experiment 2 of Tauzin et al. (2020), I observed a comparable proportion of lateral points in the LQ front conditions: the original study reports 35.71% lateral points; I found 39% across both close and far LQ front conditions, or 39% in the close LQ front condition and 45% in the far LQ front condition. The latter was more comparable to the original due to the position of the items. A more exhaustive descriptive summary is found in Table 4.1.

Table 4.1. Descriptive summary of lateral point proportion across trial types.

Trial type	**Frequency of lateral points**	**Mean proportion of lateral points**	**SD of lateral points**
Close HQ Front	29	34%	48%
Close LQ Front	32	39%	49%
Far HQ Front	15	18%	39%
Far LQ Front	38	45%	50%

I performed two planned contrasts between the HQ front and LQ front conditions, between the close and far configuration pairs. These contrasts sought to confirm the presence of a laterality effect in the close and far configurations, respectively. Family-wise type 1 error rate was set to .05, and the comparisons were performed with a Bonferroni correction for a critical p-value of 0.025 for each planned comparison.

A chi-square comparison revealed a significant difference between the far conditions, $\chi^2 = 14.58$, $p < 0.025$, $\phi_{\text{far}} = 0.30$. The far LQ front condition observed more lateral points than the far HQ front condition, which is compatible with the original laterality effect. In the second comparison, I failed to detect a laterality effect in the close configuration, which increased the communicative ambiguity by increasing physical proximity of the referents. A chi square comparison between the close HQ front and close LQ front conditions did not reveal a significant difference, $\chi^2 = 0.36$, $p = 0.55$, $\phi_{\text{close}} = 0.05$. The laterality effect in the close conditions did not reach significance (although numerically, it is in the compatible direction).

A comparison of the two laterality effect sizes (of close and far configurations) informed either the disambiguation and gravity hypotheses: in the close configuration, the disambiguation hypothesis treats the close configuration as one with more risk of ambiguity, and would assume a greater laterality effect than in the far configuration. Under the gravity hypothesis, the laterality effect is due to the difference in position between the items; a smaller change in the position in the close configurations would then yield a smaller laterality effect. The results favour the gravity hypothesis: the laterality effect was larger in the far conditions than in the close conditions ($\phi_{\text{far}} = 0.30$, $\phi_{\text{close}} = 0.05$)

Figure 4.3. Frequency of lateral points per condition.

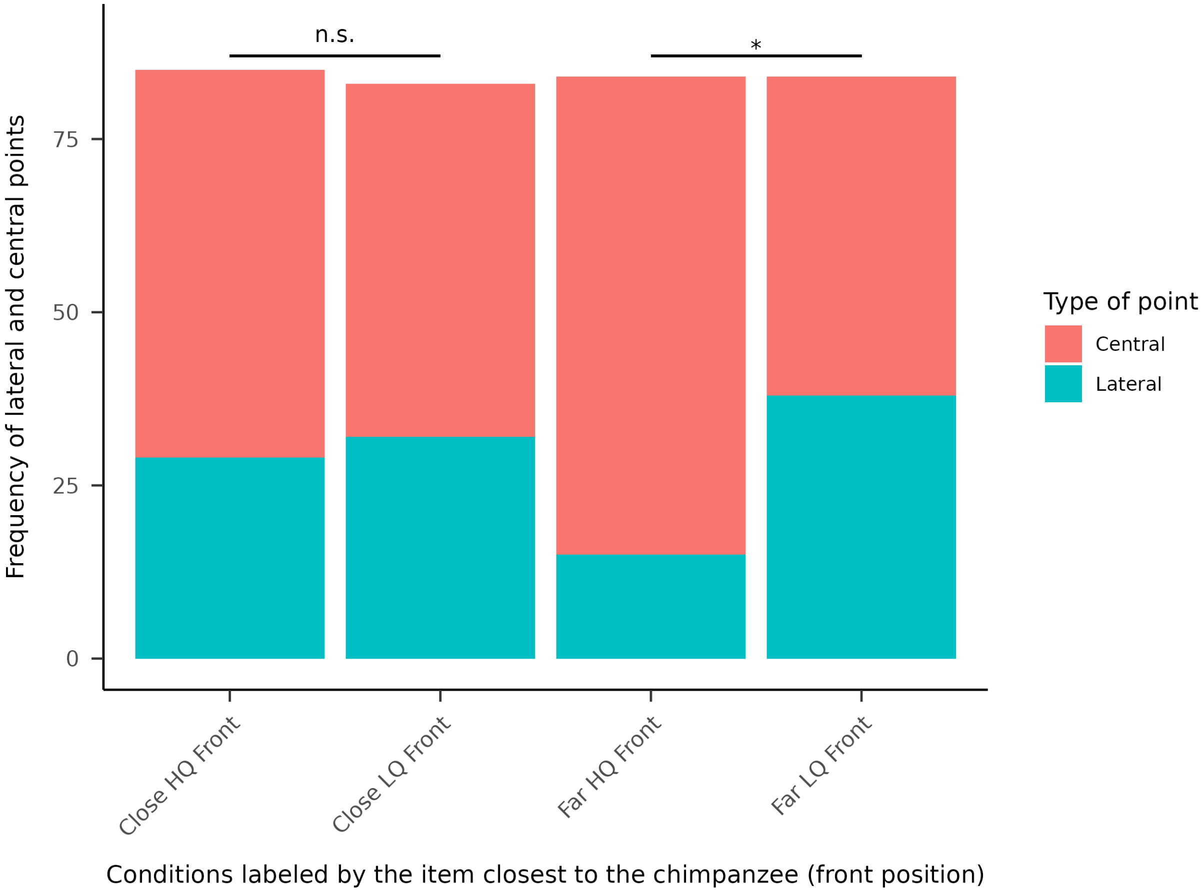

Note. Horizontal lines denote planned contrasts reported above. Significance levels set to p = 0.025.

Inspection of the descriptive scores motivated a further exploratory analysis. When frequencies of lateral points of each condition are arranged in ascending quantity, the conditions are ordered in the same way if they would be according to the distance between the mesh panel and the HQ referent. Meaning, the frequency of lateral points per condition increased numerically along with the increasing distance of the HQ referent from the mesh panel: far HQ front < close HQ front < close LQ front < far LQ front. The frequencies did not seem to increase linearly. Numerically, the differences between the HQ item position closer to the mesh seemed to correspond to a greater difference in lateral point frequency than in the farther position differences, suggesting a logarithmic relationship. In order to verify this, I compared two logistical regression models, both of which contained trial type (ordered in ascending distance) and trial order (across sessions) as the fixed predictor variables, subject as a random variable, and the frequency of lateral points as the binary outcome. The difference between the models was in whether the trial type predictor was log-transformed or not, which allowed me to examine whether a logarithmic function fit the data better than a linear.

Indeed, the AIC score of the logarithmic model, $\mathrm{AIC}_{\mathrm{log}} = 370.92$ was smaller than that of the logarithmic model, $\mathrm{AIC}_{\mathrm{lin}} = 372.51$, supporting the notion that the lateral points followed a logarithmic function more likely than a linear function.

4.2.3 Discussion

In experiment 1, I found that the chimpanzees modified their pointing according to the spatial arrangement of HQ and LQ referents on the table in front of them. The laterality effect, i.e., pointing laterally more often in the HQ front than the LQ front conditions, was replicated in in the far conditions, but did not obtain in the 'close' conditions. The former, but not the latter, is compatible with the notion that the chimpanzees attempted to 'point away' from the distractor and towards the target item in order to communicatively disambiguate the target referent for the interlocutor.

More likely, the position of points co-varies with the distance of the HQ referent. The HQ item may exert a 'gravitational pull' for the point's centrality: the closer it is, the more central the points are placed; this distance should also increase logarithmically so that smaller distances exert more of an influence than greater distances. The distance of the HQ referent is the main driver of lateral point frequency change, and so smaller changes between distances would amount to smaller laterality effects. A comparison of the close and far conditions' laterality effects (corresponding to smaller and larger distance changes, respectively) revealed a greater effect size in the far conditions than the close conditions, supporting the gravity hypothesis over the disambiguation hypothesis. The HQ referent position interacted with the point position non-linearly: smaller distances from the mesh may have had a greater influence on the point's position than longer distances. A log-transformed outcome variable proved to be somewhat of a better fit for a linear model compared to the non-transformed model, suggesting a logarithmic relationship.

This is broadly compatible with the overall interpretation of Tauzin et al. (2020) data. Pointing modifications (i.e., the laterality effect) were not sensitive to the interlocutor's knowledge state, and therefore were unlikely communicatively intentional in nature. Instead, the present sample of chimpanzees likely modified their pointing as a result of tracking the distance of the HQ referent; their points arranged themselves according to distance estimates, so that they pointed more laterally in conditions where the HQ item was farther, and they likely ignored the LQ item altogether. This is also compatible with the Gonseth et al. (2017) effect by which the pointing position matches the position of the referent. It does not, however, suggest disambiguation in the sense of specifying a correct referent over a distractor.

A better empirical situation to examine the relationship between point and referent position would allow for varying the position of a referent on a more fine-grained level, i.e., in small, regular increments on the sliding panel, as well as coding the position of the point in a more continuous manner, i.e., by measuring the distance of the point from the centre and bottom of the mesh.

These results complement the Tauzin et al. (2020) findings, further supporting the idea that object choice interactions with 'built-in' communicative ambiguity are not resolved by the chimpanzee by considering mental, pragmatic factors, but are likely guided by the spatial arrangement of the referents alone. Another part of experiment 2 of Tauzin et al. (2020) data pointed against the disambiguation idea: in a condition where there was no food in the front position, but only an HQ item in the back position (no food front condition), the apes showed more lateral points compared to the LQ front condition. This is similar to the laterality effect, but in the no food front condition, there was no need to disambiguate (there was no competing referent), so the source of the effect was unlikely to be pragmatic. In that setup, however, both locations were marked by saucers that could hold food items. In the no food front position, the saucer was empty. It is plausible that the empty saucer could have qualified, for the chimpanzees, as a liable candidate referent to avoid, in which case the obtained difference would be compatible with active disambiguation. Chimpanzees were reported to point towards empty plates that previously contained food (Bohn et al., 2015), lending plausibility to this explanation.

To eliminate this possibility, in experiment 2, I conducted a close replication of just the no food front vs. HQ front condition comparison *without* saucers (experiment 1 also did not use saucers). Should communicative disambiguation drive gesture modification in this setting, there should be no difference between the conditions; a central point would be equally unambiguous in either case. Alternatively, the gravity hypothesis would expect more lateral points in the no food front condition because the HQ referent is in the farther position than in the HQ front condition.

4.3 Experiment 2

4.3.1 Methods

Participants

I tested the same seven participants in the same housing conditions as in experiment 1.

Materials and Procedure

Same materials and food items were used as in experiment 1, with the exception of the cardboard barrier, which had to be replaced after the participants stole and destroyed the first one.

The overall procedure of experiment 2 was the same as in experiment 1, apart from the way I baited the sliding panel. In this experiment, there were two conditions, with positions in the 'far' positions (as in Experiment 1) at the ends of the sliding panel:

1. HQ front: an HQ item was placed closer to the mesh, and a LQ item in the farther position
2. No food front: there was no item in the front position, and an HQ item was placed at the farther position.

Figure 4.4 Graphic illustration of the experimental setup for both conditions of experiment 2.

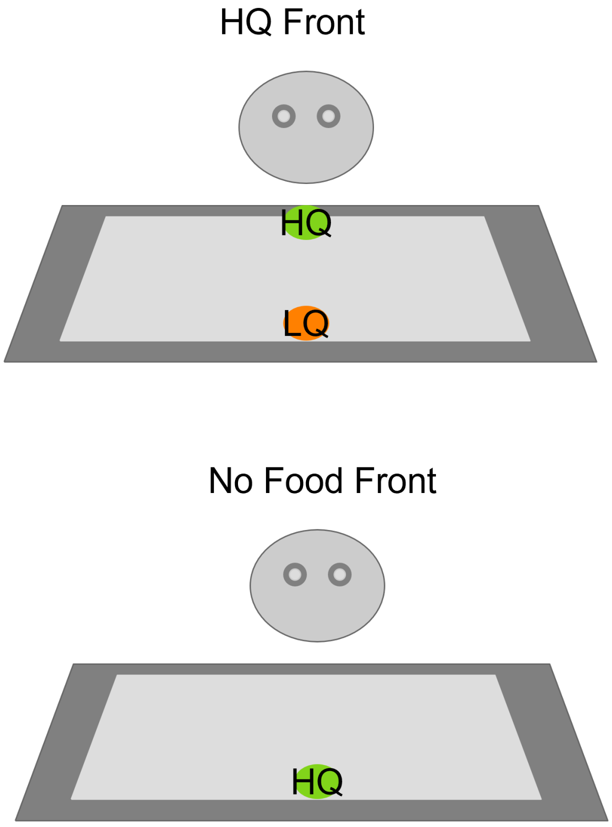

In the HQ front condition, I responded to a central point by sharing the HQ item, and to a lateral point, by sharing a LQ item. In the no food front condition, I always shared the HQ item, since it was the only item on the panel.

4.3.2 Results

I used R version 4.4.2 (2024-10-31) for all analyses.

Across both conditions, the chimpanzees pointed laterally in 39.39% trials, which in itself is comparable to the rate in experiment 1. Interestingly though, in the no food front condition, the chimpanzees pointed laterally in 59.76% trials, which is higher than in the far LQ front condition of experiment 1 (45%).

As in experiment 1, a separate coder coded 20% of the recorded session. The inter-rater reliability here was weak, $\kappa = 0.51$, $p < 0.05$.

The No food front condition induced more lateral points than the HQ front condition, $\chi^2 = 28.31$, $p < 0.05$, $\phi = 0.41$. Despite there not being a competing referent to provoke communicative ambiguity, the no food front condition yielded more lateral points than the HQ front condition. This pattern is compatible with the gravity hypothesis, which expects the point position to be governed by the HQ item position. Further, since there is no competing item at all, only the target HQ item could have directed the points' position, which can support the idea that the laterality effect is driven by considerations of only the target referent's position.

4.4 General discussion

Across two experiments, I tested whether a specific type of gesture modification is driven by communicative ambiguity, or non-social distance estimates (i.e., the gravity hypothesis). I gathered support for the latter in both experiments. In experiment 1, I pitted spatial proximity and communicative ambiguity against each other: configurations with candidate referents (food items) placed closer together on the sliding panel exerted less effect on distance estimates (according to the gravity mechanism), but more on communicative ambiguity (according to the disambiguation mechanism). The inverse was the case for the far conditions. Close conditions yielded a smaller laterality effect that did not reach significance, whereas the far conditions provoked a greater laterality effect. This is compatible with the idea that the distance of the HQ referent (target of the chimpanzees' intention) is the driver behind gesture modification and that the chimpanzees did not disambiguate the specific referent for the recipient.

Assuming distance estimation is a good way to explain gestural behaviour in this study, in experiment 1, I went on to examine the relationship between actual and perceived distance inasmuch as it is expressed in the laterality of its position on the mesh. A model with a log-transformed outcome variable (position of point) was a better fit for the data than the non-transformed model, however, the AIC score difference was small. It is plausible that, provided some methodological changes, the difference would become more prominent. For example, one could code the present outcome variable on a continuum, e.g., distance from the centre-bottom of the mesh, rather than a binary central/lateral coding scheme (as I have done here). Further, one could vary the position of the target referent with smaller, equidistant gaps between the positions of food items. This would likely afford a better empirical condition for specifying the interaction between referents' distance estimates and modification of gestures intended toward it.

Notably, the inter-rater reliability of gesture classification (central/lateral) was strong in experiment 1 ($\kappa_{\mathrm{E1}} = 0.83$), but was fairly low in experiment 2 ($\kappa_{\mathrm{E2}} = 0.51$). Plausibly, the chimpanzees' points were not unambiguously positioned enough to determine their exact position. Indeed, the chimpanzees sometimes gestured in quick open-handed motions, which was presumably hard to interpret. A discrepancy between coders could have arisen due to the viewing angle: I read and responded to their gestures live, but an independent coder used videos of the test sessions recorded from over and behind my shoulder. The angle may not have been ideal for reading the position of points.

In experiment 2, I eliminated another explanatory alternative to communicative disambiguation as a driver of the rate of lateral points. In experiment 2 of Tauzin et al. (2020), the apes pointed laterally even when there was no competing referent in the front position. They may have construed the empty saucer in the front position which normally held food as a liable candidate referent, and were thus pointing away from it. In experiment 2 of the present study, I did not use saucers, which precluded 'empty points' but I still observed a comparable effect of lateral points as in the original study, suggesting further that it was the proximity of the HQ referent that drove the position of pointing gestures, not pragmatic communicative factors.

Similar to the Tauzin et al. (2020) conclusion, I suggest that the manual gesture requests are driven by a minimal sense of context sensitivity, albeit one without consideration of the mental state of the recipient. Sensitivity to the context in the sense of spatial configuration of the referents likely satisfies the notion of 'straightforward' Brentanian intentionality, by which the chimpanzees are likely tracking the spatial position of their point's referent, but it does not reach a communicative intentionality by which the goal is to inform someone of something, and which requires maintenance of the referent embedded in the intention. It is not clear whether this would satisfy the Townsend et al. (2017) criteria of intentional communication. Presumably, the chimpanzees wanted to obtain the HQ item and therefore were holding a goal about something in the world, but it is not also clear that they engaged in maintenance of the referent for communicative purposes, i.e., disambiguated it against distractors.

A final methodological note: the task may carry a pragmatic inconsistency. The experimenter repeatedly (across trials) ‘asked’ the chimpanzees to express their choice, and the logic of the task relied on the chimpanzees’ points being informative – i.e., what they wanted. However, their preference was already established in pretest preference tests, as well as across trials, so the experimenter could have known what they wanted already. If the chimpanzees were insensitive to the situation’s pragmatic factors, the test might not be appropriate to tease out their communicative abilities. On the other hand, if they were sensitive to the pragmatic factors, they would not necessarily be motivated to express redundant information. It is not clear to me how to avoid this methodological issue, but I will reflect on it in the general conclusion.

Chapter 5. Species-typical constraints of chimpanzee cognition: leveraging two-system accounts

5.1 Introduction

The introduction of the thesis stopped the review of nonhuman primate ToM literature with a brief reflection on several two-system approaches, briefly weighing them against the available evidence. The empirical chapters explored different operational approaches to teasing out more testable situations from two-system approaches. Roughly, the methodology followed an implicit approach of empirical psychology which compares theoretical models on grounds of how well they are corroborated by evidence, and how well they can be instantiated empirically. This competitive corroboration approach is not the only theory appraisal route, though. Gigerenzer (2017), for example, outlines a 'theory integration approach' by which seemingly isolated theoretical constructs can be analysed, integrated, and dissociated on operational and conceptual grounds. This allows for unifying conceptual and operational constructs on grounds which seemed dissociated for lack of a thorough analysis, or dissociating constructs that were treated nominally equivalent but dissociate conceptually or operationally. Effectively, this is an attempt to make meaningful distinctions with respect to the disciplinary empirical and theoretical commitments that psychologists engage in. Ideally, this would lessen problems of too isolated theoretical and empirical efforts which do not inform one another. Mischel (2008) recalls a witticism that may come to be known as the toothbrush problem: 'Psychologists treat other peoples' theories like toothbrushes – no self-respecting person wants to use anyone else's. This is a problem in that it renders all empirical corroboration and conceptual analysis local to those efforts, and does not build a more comprehensive epistemic network.

A theory integration approach commits psychological constructs to their conceptual relations and operational implications. Only the aspects which can be conceptually and operationally defined, and put into relation with other theories and the data are regarded as relevant. Apart from the conceptual and operational levels, I suggest expanding the methodology to also include a third level of analysis: the fundamental assumptions, or the 'hard cores' of theoretical programs. These are commitments about the basic phenomena that a theory purports to explain that are shared by all of its conceptual and operational implications. For example, a representational cognitive theory will assume that the mind's internal states relate to, or explain, external referents. We may contrast this to behaviourist theories which do not assume mental content, or a processing mechanism, but trace the reinforcement history that shaped it. The fundamental assumptions of a theoretical program determine not only how a phenomenon is interpreted, but what the phenomenon is in the first place:

an internal representation, or a reinforcement plan? A cognitivist and behaviourist theory will not share their fundamental assumptions, and they will be orthogonal as to a) how they construe basic phenomena, b) conceptual and c) operational implications. A protocol that informs a cognitivist psychological construct, like minimal theory of mind, need not contain operational commitments to reinforcement history. Inversely, two cognitivist models can ‘share’ a protocol so that the outcomes inform both, e.g., so that the data is compatible with one model but precludes the other. When two researchers share fundamental assumptions but fail to relate their theoretical constructs to one another’s, they miss out on this form of ‘local’ empirical corroboration. In this chapter, I will attempt to pick up where the introductory chapter left the review of theory of mind theoretical accounts and ‘run through’ what I see as relevant points at foundational, conceptual, and operational levels of analysis of interesting two-systems theory of mind models.

I imagine a workable epistemic sketch as follows: the fundamental assumptions determine the basic ‘object’ of inquiry, e.g., conceiving the mind as representational. These assumptions are ‘upstream’ of conceptual and operational definitions. A conceptual definition gives a description of the causal structure in which some relevant phenomenon is implied, e.g., a specification of the proprietary input or the transformational rule that generates the output. An operational definition might apply this to a particular protocol and determine measurements, effect sizes, etc. For the kind of cognitive theories of mind that I worked with throughout this project and thesis, we may talk about mental representations or intentional attitudes about, or directed at the world as basic objects, whose character will inform domain-general and domain-specific constraints that characterise the representational dynamic. Further downstream, the operational definitions will determine an animal’s behaviour response in protocols like anticipatory looking in a false belief task, head rotation in a gaze following test, etc.

Motivated by two-system theories of the theory of mind (e.g., Apperly & Butterfill, 2009; Nagel, 2017) remarked upon in earlier chapters, I applied the heuristic of dividing the conceptual problem space of two types of representations – primarily internal or mental, and primarily observational or physical. Thus far, I examined different operational expressions of this heuristic in the domains of object representation and communication. In this chapter, I will ‘zoom out’ to some domain-general constraints of the mind, and then return to conceptual and operational descriptions of other domains of thought as proof of concept: object individuation, inference, and communication.

5.2 The basic object of mental representation

At the outset, one should outline the initial commitments of what it means to think about something or to hold a representation of something in the mind. At the most basic level, we can imagine an animal to hold an internal (mental) state which relates to something in the world. The relation between the internal state and the referent in the world cannot be a simple correlation or concurrence; the representational function must somehow capture the relevant 'aboutness' character of the relation. As a clue, take Langer's (1957) instructive distinction between signs and symbols: signs are correlates or concurrences between states in the world. They can be natural, spontaneously occurring relations, so that a patter on the roof is a sign of rain, or they can be artificial so that a ringing of a bell announces dinnertime. In both cases, signs announce something in the world, so that an animal that recognises a sign can direct their attention and action appropriately. She need not form any internal content describing or explaining the referent. On the other hand, a symbolic relation does not simply direct attention or action to the world; instead it invokes some internal mental content. For Langer, the crucial distinction is that a symbol directs to an internal mental state, while a sign guides action in the physical world. I take this broadly construed function of a symbolic or representational relationship is inherent to the representational explanation of the mind.

Perner (1991) offers a more tractable conceptual description: for a state to represent something in the world, it must represent a referent as something, or in a particular way. As Langer (1957) points out, it is not enough for a representation to co-occur or extend from what it represents, just as it should not be enough for it to copy it (if that were possible). It must somehow reinterpret the referent into something recognisable to the mind. This yields a dissociation between what is being represented – the referent, and the way specific way it is represented in the mind – representation. The use of 'representation' here should not be mistaken with the representational relation or function, which describes the relation between the (more narrowly construed) representation 'in the mind' and the referent 'in the world'. I use the term in the narrow sense, to denote a state in mind that is in a representational relation to the referent, which is 'out there' in the observational, physical world.

Further markers of the representational function can be briefly summarised as follows: a representation is asymmetrical to the referent so that their relation is directional. A representation is about the world, but the world is not about the representation. If it individuates a part of the world, like an object, a representation can be about this object alone, and not also include other, identical objects. There should also be a way to misrepresent the referent, so that by virtue of reinterpreting the referent as some representation, it gains the opportunity of misconstruing it, or more broadly, being incorrect about it. Relatedly, a representation can be about something that does not exist in the world – either because it is false, like a misrepresentation, or because it depicts something that does not or cannot exist, like an Escher's drawing. Perner's (1991) markers are relatively formal ways of describing the representational relation, or function, and would deserve more elaboration. Due to practical constraints and in the spirit of the chapters' aims, however, I choose to build the formulation of a representation in psychological terms. Mental representation may be a philosophical kind, but in so far as it is interesting to a psychologist, it should also embody a theoretical and operational psychological form.

For a psychologist, it is important to show that a representation can 'go beyond' an observational, perceptual experience in a way that it cannot be reduced to observing observational contingencies. A representation of a chair is more than the sum of some observations: it is determined by a representational function that unifies those experiences under a certain chair-esque principle. The question of where such a function comes from is an open, possibly intractable one. The standard cognitivist solution is to posit some innate prespecification of at least basic psychological entities on top of which animals form more complex forms of knowledge. This specification, in turn, is not usually explained, which is a clue to the fact that it might be the 'bottom line' or the hard core of the representational explanatory strategy.

Likely, there should be some causal relationship between the representation and the referent, e.g., innervation at the sense organ which can presumably lead to a mental, representational experience. However, in all but possibly the most basic neuronal activation, this relationship is opaque, or intractable. It should be possible to allow some presumption of a causal relation without also specifying its nature, just as it is possible to suppose that the mind is a product of natural selection without knowing an exhaustive (or even accurate) causal history. I will refer to the disconnect between the observational referent and the mental representation as the inductive gap, i.e., the opaque representational function which takes observational (or similar) information and generates related, but not straightforwardly extensional, internally specified mental representations.

The induction gap, thus construed, opens up the space for representation to exist without maintenance through relation to observational information. We can think about the world without ongoing experiences of it, but more so, we can be incorrect about the world, or even in terms which *cannot* be expressed observationally at all. The most obvious cases are concepts pertaining to non-observable entities, e.g., atoms, sub-atomic particles, temperature, energy, personality traits, mental representations, etc. These are real entities, in the sense that they can be meaningfully thought about, measured, and related to other similar entities, but crucially, they need not be observable themselves. Scientific and linguistic statements like 'chimpanzees have a ToM', 'Cesar crossed the Rubicon', or 'all crows are black', etc., cannot conclusively be reduced to observational terms: an observation that would conclusively and unambiguously determine these types of propositions does not exist. Either the observational space, i.e., ‘all crows’ is impossible to build, or the evidentiary event cannot be replicated in the present. The underspecification of the inductive gap, or the representational function, opens up the space for misrepresentation, and furthermore, nonexistence of mental content. Any attempt to reduce the representation to its observational referent will, presumably, encounter some issue of intractability.

At least some cognitive entities and relations should interact without the friction of induction, though. These find expression in formal logic and mathematics, where reference to observation is not at all necessary for maintenance of content, and it is instead upheld by relations to other nonexistent entities. Deductive logic enjoys a privilege, much envied by empirical science, of absolute certainty, like in statements like ‘all crows are black’. The validity is not conferred by observational validation, but through consistency with other statements, also taken to be true. The ALL operator, here, enables a more radical decoupling from observational reality, in which the primary mode relates mental contents to one another, rather than to more basic observational information.

Thus, the mark of a mental representation is in its decoupling with the physical or observational world so that it at least cannot be straightforwardly related to it. The opaque function of this decoupling I will refer to as the inductive gap. Presumably, there are degrees of this decoupling, such that some mental representations purport (but ultimately fail) to relate to physical entities inductively, and some are engaged in relations with other mental entities by means of deductive operations. Such is my cursory reading of the basic object of a cognitivist, or representationalist explanation.

5.3 Domain general models of the mind

Representations are *about* the (observational) world in that they contain information about how to organise or interpret it. Plausibly, this at least primed philosophical proposals which conceive cognition as a propositional entity, somehow similar to language. More explicitly, Fodor (Fodor, 1975, 2008; Fodor & Pylyshyn, 1988) maintains that language is the best model for the mind because it is the only example of a system organised in discrete parts and operated upon with some combinatory rules. A propositional, language-like system, is the best way to model the mind – so goes the basic argument for language of thought (LoT) from systematicity.

LoT was originally proposed as underlying all possible cognition, and not only that of language-able humans; its systematic relations do not depend on natural language competence or explicit symbolic systems. Rather, systematicity here refers to operational regularities like transitivity, symmetry, hierarchical structure, etc. Regardless, comparative cognition researchers usually work against the prior of animals being non-cognitive, or at least not in a human-typical way (e.g., Povinelli & Vonk, 2003). Plausibly, though, barring animals from cognitivist explanations is throwing the baby out with the bathwater, and runs counter to the initial commitments of LoT (or similar accounts) to explain the mind per se, not just the human mind. animals, too.

Against the background of this sceptical outlook towards animal minds, Camp (2007) re-examined the LoT argument to justify an alternative model which is still compatible with the systematicity argument. The original LoT argument took a language-like system to be the *only* type of system capable of meeting systematicity, i.e., accommodating recurring discrete parts combined according to some combinatory rules. Camp reformulates the argument into a 'weaker' form according to which propositional language is one, but not the only model that meets this demand. She suggests an alternative 'cartographic' representational format according to which representational information is not organised, e.g., in a syntactic manner, but instead adheres to physical constraints of objects and entities, so that the way a mental 'map' is construed will resemble the physical arrangement of the world. Maps, here, are taken to be visual-esque and resemble ordinance survey maps, satellite images, or pictures of the external world, or rather, they will share some of their representational constraints. Maps, as mental representations, are more sophisticated than simple pictures and allow for symbolic depiction by way of icons and shapes representing objects in physical space. Importantly, for Camp (2007), a cartographic representation must follow the principle of isomorphism, by which the depiction on the map must somehow resemble the physical arrangement of things in the observational world. In contrast, language-like representations stand in an arbitrary relation to what they represent, so that a word or a sentence need not resemble its observational referent. Maps, however, are isomorphic with the observational world, and they rearranged in a way similar to how objects in the observational world are, too. A river flowing through a valley should be depicted as something like an undulating line, that respects both its shape and its relation to other geographic features.

A representational format must have a set of proprietary constraints that govern its expression. For Camp, isomorphism is the main overarching constraint for cartographic representation. While I agree that it is likely true that maps will need to correlate with worldly physicalities, I think the isomorphism principle fails to show a *proprietary* representational constraint. It delegates the description of its representational nature to the physical environment, which, presumably, it faithfully copies into an internal representation. Regardless of how the relationship between the physical world and the representational structure is established, the latter must be able to determine its own, proprietary constraints, not by delegating the work of specification outside its own means. Representational constraints must come from 'within' the description of the format itself, not outside it.

In order to build a set of proprietary constraints for the cartographic representational format, we may look for precedence in other formats like the LoT. Quilty-Dunn et al. (2022) offer a proprietary set of constraints for the LoT representational format. In the following section, I will elaborate on these properties, and use them as a frame of reference for their cartographic counterparts.

5.3.1 Propositional representational format

Qulty-Dunn et al. (2022) developed six discrete empirically tractable constraints of the LoT system. LoT should exhibit 1) discrete constituents, 2) role-filler-independence, 3) predicate-argument structure, 4) the use of logical operators, 5) inferential promiscuity, and 6) the handling of abstract conceptual content. I organise these properties (and add one) in four principles that should exhaustively describe the system in its: 1) combinatory rules, or how its constituents are organised, 2) operational rules, or how information can be manipulated and transformed, 3) content depiction, or how abstract or concrete the represented objects are, and finally, 4) referential structure, whereby incompatible and counterfactual content is either possible or not. These discrete constraints can be helpful both for theoretical analysis and empirical instantiation, but it should be possible to regard the set of constraints as an expression of a singular LoT-esque natural kind. It seems entirely plausible (and, to me, likely) that these are interrelated aspects of a unitary psychological construct. Presently, however, the goal is to provide a methodologically instructive outline that lends itself to theoretical comparison, as well as operational definition.

Syntactic combinatory principle

Just like syntax determines how natural language is organised, so does LoT exhibit syntax-like combinatory rules. The first three of Quilty-Dunn et al. (2022) properties correspond to the combinatory organisation of constituents. To this end, they outline three principles: 1) discrete constituency, 2) role filler independence, and 3) predicative argumentative structure. First, discrete constituents mean that constituent parts have individual and separable features, and removing single parts from a whole representation does not lead to a collapse of the whole structure. If, in a representation 'A coffee cup is on the table', we remove the unit 'coffee' *a* structure persists, though its meaning may change. Other parts of the structure are not interdependent so the whole structure breaks down if one part is altered. Second, the parts in a Lot structure are organised according to a functional role which is independent of the depicted content. In other words, the syntactic role which determines how a unit is organised is independent of what information that unit carries, so that in 'Mary loves John' and 'Sue loves John', the structure of the expression is held constant, but the depicted meaning varies. The organisation of the constituents is independent of the content it depicts. Once these two properties are in place, the ingredients for the syntactical predicate-argument structure are in place. Discrete constituents are organised according to some functional role inherent to the representation. This means that a structure cannot be replaced with an unstructured whole, for example, ‘John smokes’ is a representation which is more than just a concatenation of ‘John’ and ‘smoke’; the syntactical predicate-argument structure is inalienable from the representation itself. Thus, these three properties outline what is for LoT analogous to what syntax is for natural language: an organisational principle that structures discrete units of information according to some functional role.

Deductive logical operations

The next two properties, under which LoT exhibits 4) use of logical operators and 5) inferential promiscuity, pertain to the operational rules for transforming content, or inference. Logical operators like IF, NOT, OR, etc. enable the expression of conditionals, negation, and counterfactuality. These operations are inherent to the LoT, and need not be articulated through external symbolic scaffolding, like written or verbal language. Instead, they can ‘proceed’ on their own accord, spontaneously or automatically. This is not an operational commitment to exact processing time, or to conscious access to the process; rather, it pertains to the fact that processing can proceed without necessarily requiring effort for locution. Of course, deductive inference is not always automatic or spontaneous. Formal mathematical or logical problems can be notoriously difficult for adult humans – e.g., see the classical Wason selection task, (Wason, 1968). Rather, promiscuity refers to the fact that logical relations are immediately implicit and readily expandable, although presumably, in practice, this is subject to performance limitations. Logical relations and inferential promiscuity are means of articulation and locution of logical propositional thought; they determine how thinking proceeds, and that it can do so spontaneously, or of its own accord.

Abstract content

Quilty-Dunn’s last property concerns the nature of the content a LoT can support, or how the represented objects fare on the dimension of abstraction. Abstraction is not an idea to take for granted, so to constrain the matter here, take it as a degree to which it is bound to perceptual information. If an object, or any content, is represented with information about how it might present to the sensory and perceptual apparatus, then it is concrete; if it is not specified this way, it is abstract. I take this as roughly equivalent to saying that concrete content is expressed in observational terms, while abstract content is not. LoT will be uniquely able to represent abstract (i.e., non-observational) cognitive content which need not readily manifest in observational terms or conform to constraints of a certain modality. Propositions like ‘the cat is on the mat’ do not explicitly designate modality-specific information like shape, texture, taste, etc., but a general membership that need not straightforwardly translate to physical, observational terms. In a sense, the information is disembodied. This is even more obvious in statements like ‘Socrates is wise’ or ‘chimpanzees have a ToM’ which are laboriously related to any observations, but are still meaningful as propositions.

Intensional reference

The final type of constraints is not described in Quilty-Dunn et al.’s (2022) list but appears as a criterion of propositional belief states in some models, at least implicitly (Apperly & Butterfill, 2009; Butterfill & Apperly, 2013; Lurz, 2009; Nagel, 2017; Rakoczy, 2012; Tomasello, 2018). It concerns the nature of incommensurability of reference, by which two representations of an objects can simultaneously share a referent, but remain incompatible with one another. Say I refer to a colleague’s new garment as a skirt. She may protest and say that the garment is a kilt. Our referential situation is said to be intensional. Mine and my colleague’s perspectives are incompatible, despite sharing a referent: there must be some concession made for a perspective’s unique vantage on something in the world. This incompatibility is evident in cases of confronting the perspective of another or tracking others’ false beliefs. For example, Sally and Anne may look down on a box placed on the floor between them and can be said to have different perceptual information based on their viewpoint, as well as different sets of dispositions with which they might approach the situation. Their mental states are rendered even more obviously incompatible if their relevant experiences differ: if the box contains an object that only one of them saw baited, there will be specific points of incompatibility between their perspectives. Regardless, since they are attending to the same object in space, they can be said to be converging on it, despite their different vantage points. This kind of intensional structure exposes a peculiar way in which perspectival mental states refer to the world: they are incommensurable with each other and not exhaustively described by the ‘objective’ or shared physical reference.

Intensional situations differ from extensional, which lack the additional perspectival, incommensurable aspect, and can be exhaustively described in observable, physical terms. If my friend and I agree on our terms and I admit that she is wearing a skirt, but another colleague’s garment is a skirt, then our labels are no longer incommensurate, and can be reduced to their observational terms. In an extensional context, there is no ‘room’ for incommensurate properties. Everything must fit ‘on the same plane’, or else can not be expressed at all. As an example of a minimal contrast between intensional and extensional contexts, take the alternative naming game (Doherty & Perner, 1998) where an experimenter and a child take turns giving names or descriptions to an object, and they cannot re-use terms. In an intensional version, the object is a rabbit toy, and the players can take turns naming it ‘bunny’ and ‘rabbit’, etc. The two names are incommensurate with each other because they both exhaustively characterise the object, and do not ‘leave room’ for each other. In an alternative, extensional version, the players take turns naming the colours on a flower, e.g., ‘green’, ‘red’, etc. In this case, the labels designate *parts* of the flower, not the whole, and therefore they leave room for each other. The colour labels are, in this case, co-extensive. If we were to take turns naming the object ‘flower’, ‘rose’, etc., we would be back in an intensional situation because these labels describe the whole object, and do not leave room for another part of the flower to carry a different name. Call these references that refer to an object exhaustively ‘holistic’. When two holistic references converge on the same referent, they do so in spite of their incommensurability. Extensional reference, in contrast, subsets a part of the observational world and allows for co-extensive referents.

In scientific discourse, as well as in common parlance, we routinely accommodate perspectival intensional situations, e.g., with phrases like ‘thinks that’, and ‘sees as’, which seem to carry the appropriate referential function. One can trivially say that ‘Matt thinks that the ball is red, but Ava sees it as maroon.’ and be able to appreciate that their unique perspectives are not readily explained by reference to the ball in physical space. It may be that propositional structures are fit for accommodating intensional references, and that LoT is an appropriate representational structure for this task.

As an interim summary before moving on to the cartographic non-propositional representational format: LoT is a theoretically tractable representational format that is organised syntactically, can perform logical operations, can hold abstract and domain-general content, and can amount to intensionally structured referential structure. Borrowing the methodology of breaking the system down to tractable principles, I will proceed with the same methodology to describe a set of non-propositional, cartographic representational format.

5.3.2 Nonpropositional representational format

To reiterate, the cartographic representational system is a non-propositional representational format that should resemble constraints of Ordinance survey maps, satellite images, and pictures of the external world. It permits the use of relatively abstract icons, shapes, and colours to designate entities, their position, appearance, and shapes in the world. It should not use linguistic symbols, phrases and sentences, though, and limit itself to icons rather than arbitrary symbols. For Camp (2007), the overarching principle of cartographic representation is isomorphism, by which maps must, to some extent, resemble what they represent. For example, a blue line that undulates across a map will resemble the way a physical river undulates around obstacles in the external world. It seems obvious and tempting, then, to say that the representation of a river on the map should resemble the actual river, or that they are isomorphic. However, this also fails to specify the proprietary constraint of the format like the above LoT analysis. It is possible, though, to apply the same four-tiered analysis to the cartographic system. Thus, I will contrast the cartographic proprietary constraints with the LoT constraints under the same four categories: 1) combinatory principle, 2) operational processes, 3) content depiction, and 4) referential structure.

Topological combinatory principle

LoT's combinatory principle resembles syntax of natural language. This means that it is structured in discrete constituents and that there is a functional structure that its parts are organised within. This structure is independent of the content the representation depicts. The organisation of a cartographic representation, on the other hand, will have a more direct relationship with the content it conveys. For example, the structure of a sentence 'the chair is next to the table' does not mirror the physical situation, but a map or a picture of a chair next to a table will place the objects next to each other, following the way they present physically. This is what Camp (2007) understood under the overarching principle of isomorphism – a correlation between the real, physical situation and its depiction. However, isomorphism, in that formulation, does not offer a proprietary constraint for how representation works, but it delegates explaining its rules to outside itself.

Following LoT's combinatory principles, we may weigh maps against the same criteria. A LoT representation is structured in discrete constituents, which means that constituents can be taken out of the representation without collapsing the whole structure. When a word is taken out of a sentence, other constituents need not be deleted. Similar can be said for a map – taking out a constituent like an icon designating the location of a bothy does not collapse the rest of the map. This does not mean that the constituents are wholly independent, though: if a house on a map is shown next to a lake, it can be thought of as a lake house. The relation between the constituents can be a part of their meaning. This goes for propositions, too – the meaning of 'cup' may change depending on whether it is in a formulation 'a green cup' or 'a red cup'. Thus, both maps and LoT exhibit discrete constituents.

LoT's role-filler independence structure, however, should set itself apart from maps. For the latter, the content of constituents is not independent from the structure in which it is situated in. In, e.g., depicting a river running through some terrain with a blue line, the position and shape of the blue line are relevant to what it designates, i.e., the shape and location of the river. This implies a dependence between the structural principle and the content it depicts: the squiggle on the map is related to the undulation of the river in space. Call this the topological principle – relations in which constituents are depicted on the map will bear on the depicted content. Putting icons on the map either closer or farther apart is not interchangeable with respect to what the map will denote. This is not true for propositions – there need not be any straightforward causal (or analogous) relation between the designated content and the functional role in which they are organised.

Thus, where LoT is structured with discrete constituents and role-filler independence, maps will also be structured with discrete constituents, but the way they are organised will load on the meaning they depict. LoT is thus organised syntactically, and a cartographic representation, topologically.

Inductive logical operations

Regarding manipulation and transformation of content, LoT deployed the use of logical operators, which is a staple of deductive reasoning. This allows it to accommodate formulations like the disjunctive syllogism: 'either p or q, not q; therefore p; therefore p'. This is a type of epistemic opportunity where a conclusion does not rely on imperfect inductive corroboration but on coherence with premises taken to be true, from which it inherits absolute certainty. This is exhibited by propositional logical systems, and more specifically afforded by logical operators like IF, OR, etc., as well as certain quantifiers like ALL, SOME, etc.

Conversely, induction does not enjoy the same degree of certainty as deduction but relies on partial data-driven corroboration. In line with the fact that cartographic representations tend towards observational terms, they are also constrained to inductive operations that rely on empirical corroboration. Here, I take inductive inference to be the kind of operation which takes as input observational information. I do not use induction in the more classical sense of updating propositional statements with observational data – rather, I leverage the contrast between empirical and deductive operations as a limiting case of a type of psychological operation. Induction represents information in terms of how it can be found in the world – in quantities, shapes, sizes, appearances, etc. This limitation does not commit the cartographic system to a particular mechanism of representing quantity, size, etc., but it prevents it from going 'beyond' this type of information onto deductive space where they can deploy logical operators and quantifiers such as IF, OR, ALL, etc.

One empirically tractable limitation that comes out of this is that a cartographic system will not approach situations with the relative quantity or probability of outcomes by representing information in terms of complete/infinite quantities – i.e., 'ALL x are y' or absolute certainties 'x must be y'. Instead, the cartographic inductive system is limited in representing some quantitative estimates, and will not generalise to judgments about totality like 'all crows are black'. Despite gathering a growing number of black crows, it will not resort to judgments about ALL crows. Rather, the sample quantity estimate will continue to rise according to some function. Similar goes for probability estimates: relative estimates are fine, but certainties should be reserved for the propositional, deductive space. Thus, where propositional deduction has certainty and expansiveness – e.g., ‘all crows are black’ – cartographic induction has probability and quantity estimates – e.g., ‘some crows may be black’.

Another possible point of dissociation between the representational systems concerns representing negation and absence. There are at least two interesting kinds of negation relevant here. For example ‘there are no pencils in the case’ or ‘a self-correcting pencil does not exist’ imply a different negation operation. The first conveys absence, but the latter conveys positive non-existence. Difference being, it is only in the former that we can find reference to some absent content, whereas, in the latter, there is no such reference. Gomez (2009) discusses these kinds of cases as pertaining to inexistence and nonexistence. Inexistence relies on similar structures like imagined physical relations in space (e.g., manual lines of point), except that the referent, here, is also imagined. In nonexistent reference, there is no such imagined referent. In the former, the referent is merely suspended but still is thought of in some inductive, observational way. It is not a true, harder case of negation, as in nonexistence.

Concrete content

The topological property determines the organisation of a map, which then stands to be populated by the semantic element. The constraint which distinguishes LoT and cartographic representation is the level of abstraction in which some content is depicted. Abstraction, here, concerns whether or not, or to what extent, the representation carries observational information, i.e., that pertaining to a particular sensory modality. For example, the proposition ‘A cat is sitting on the mat’ does not subset cats or mats of particular appearance; it states general membership of cats and mats, and need not immediately invoke a visual depiction. On the other hand, a cartographic or pictorial representation of a cat should be descriptive of at least some of its features like fur colour, size, etc. Propositions, conversely, can persist without a mental picture. This is further supported by observing more extreme cases of propositions without an obvious observational correspondent. Propositional statements like ‘All swans are white’ cannot be exhaustively described in observational terms because the space they describe is infinitely large – they talk about *all* swans, not a determined finite sample of swans. The same goes for quantifiers like ‘some’: there is no way to reduce ‘some swans’ to any particular finite sample – ‘some’ deliberately omits a quantitative estimate which would make it easy to depict on a map by placing n of icons in an area. Other statements like ‘chimpanzees have a ToM’ are unclear as to what kind of physical or observable state of affairs they would refer to; reducing such a proposition to observational terms proves laborious and contentious.

Cartographic representations should contain some physical or observational information in their depiction. Even when icons are of an abstract form, they will have a non-arbitrary spatial relation in the map. Thus, apart from the structure being connected non-arbitrarily to the content it depicts (topological organisation), the constituents will also carry some observational, domain-specific specification, which is what I take to mean that their content is concrete. Propositions, on the other hand, are uniquely able to carry general, abstract information about what they depict.

Extensional reference

Intensional referential situations are those that involve perspectival and counterfactual information which can stand incommensurate with each other despite being directed at the same referent. Extensional situations, inversely, do not accommodate incommensurate but converging referential structures. There is no ‘room’ for such conflict in extensional contexts, which can only depict how things look from a single perspective – or rather, there are no perspectives since there is no opportunity for uniquely perspectival content.

When there are no alternative perspectives or possibilities for an animal to consider, we may say that her knowledge describes her world totally, or holistically. Each constituent statement or observation, likewise, is a part of the animal’s holistic knowledge about the world and is not regarded in isolation. In Quine’s formulation ‘our statements about the external world face the tribunal of sense experience not individually but only as a corporate body’ (Quine, 1951, p. p19) The corporate body is the holistic ‘picture’ of the world which sets the attitude to consider the totality of a situation. This is not to say that a cartographic, holistic representation includes all possible or necessary information about the world explicitly, but implicitly, which is to say that whatever information that could be included in the representation is part of the same holistic knowledge.

Every psychological agent has a holistic picture that runs down to their observations. Likewise, two vantages, no matter how similar, will always differ in some intractable way. A propositional solution is to ignore the holistic and isolate vantages only in so far as they relate to particular referents. They converge not by juxtaposing incommensurate holistic world pictures, but by isolating agents’ relations to only some constituents: ‘Sally and Anne are both looking at the box.’ Surely, Both Sally and Anne have their own holistic worldviews, but a propositional formulation can ignore that in favour of considering only their relations towards the box. It can do that intensionally, by opening up unique vantages to a shared referent.

Without an intensional structure, a cartographic representation cannot make constituents particular in this way. Instead, a cartographic animal can track an agent’s perceptual affordance as part of her own holistic world. She can track their perceptual access to her holistic knowledge (world picture) but can leave no opportunity for that agent's knowledge to be incompatible with her own. Knowledge states, in this sense, cannot be false, they are only marked by successful or failed perceptual access to a holistic, factual body of personal knowledge. If perceptual access is granted, then an agent will partake in some of the personal knowledge. If not, she will be barred from it.

Several theoretical analyses already made the important distinctions of knowledge and belief, or factive and nonfactive epistemic content, to distinguish types of theories of mind (Nagel, 2017; J. Phillips et al., 2021; J. Phillips & Norby, 2019). The remarks in this section are compatible with these, but are more domain-general constrains, and are not contained to a particular domain or modality. For example, similar should apply to considerations of mutually exclusive events, such as in future planning. No true incompatibility with the holistic worldview is possible, and no true alternative event is possible.

Summary of constraints

In sum, there are four main types of constraints for any representational format: the organisational principle, the mode of inferential operations, the abstractedness of content, and the referential structure. Propositional systems like LoT are organised syntactically, are uniquely capable of deductive inference, handling abstract content, and intensional structures. Nonpropositional systems like the cartographic systems are organised topologically, are limited to inductive inference, concrete content, and extensional reference. These properties are summarised in Table 1.

Table 1. Summary of constraints according to type.

General properties	**Propositional LoT**	**Nonpropositional maps**
Combinatory principle	Syntactic: discrete constituents; role filler independence	Topological: discrete constituents; role filler dependence
Operational rules	Deductive: logical operators; inferential promiscuity	Inductive: quantity estimates
Content	Abstract: amodal information	Concrete: modality specific information
Reference	Intensional	Extensional

5.4 Empirical proof of concept

In order for this theoretical analysis to work as a good scientific model, the meaningful distinctions made at the conceptual level should be relevant on the operational level, too. The dichotomous picture of cartographic and LoT systems should manifest operationally as signature limits in performance, or better yet, format-typical performance patterns. Should that be the case, then the analysis will be further vindicated. For practical reasons, I narrowed the scope to cover only nonhuman primate data, even though a more complete analysis would include at least human children and other animal species. Further, there is no in-principle reason to constrain the discussion to only some cognitive domains, but again, for practical reasons, I review only theory of mind, object individuation, inference, and communication. The review is a proof of concept and not a comprehensive, conclusive overview. Some of the relevant data has already been reviewed in previous chapters, and for some sections, I will gloss over empirical aspects and focus on the theoretical distinctions alone.

5.4.1 Theory of mind

In the introductory chapter, I reviewed the data concerning nonhuman primate ToM data, particularly that which can deliminate representational structures found in similar operational situations. Briefly, there is plenty of evidence suggesting that nonhuman primates reportedly follow gaze and represent knowledge-like states. They also fail to act appropriately in several live belief-involving tasks but succeed in some spontaneous behaviour-based tasks.

Since the empirical side is already established, in the following section, I will pick up where the theoretical discussion stopped and proceed with a more detailed analysis of the theoretical constructs and how they may or may not map onto the propositional/cartographic problem space.

Minimal mindreading

In Butterfill and Apperly's (2009; Butterfill & Apperly, 2013, henceforth A&B) model, the minimal mindreader represents registrations through certain conceptually tractable principles. Briefly, registrations are constituted by a relation between the agent, what is encountered (e.g., an object), and the location in which it was last encountered. These registrations will pertain to agents' actions in the world, so that correct registrations will lead to successful actions, and in case of incorrect registrations, the agent is presumed to act to the best of her knowledge about the location of objects. Thus, in a false belief task, when the agent fails to witness the object changing locations, the minimal mindreader will assume her to act to the best of her knowledge, according to her last successful encounter at the old location. In this case, registrations converge with the prediction that a propositional theory of mind would also make, so they can be said to emulate beliefs, at least operationally. A&B's principles (encountering, registration tracking, etc.) are not necessarily psychologically instantiated in the same way that the cartographic and LoT properties are, so they could stand to be examined in relation.

Agents' encounters with objects or entities in their perceptual field are expressed in their physical relations: the set of objects comprising the perceptual field is constrained by the affordances of what the agent might see (without necessarily invoking the mentalising implications of seeing). Objects which are occluded, not illuminated, or positioned far outside an agent's direction of gaze will not lead to successful registrations, whereas those likely to afford sensory or perceptual detection will. Thus outlined, perceptual encounters resemble how cartographic perceptual works, too. Access, here, takes the form of a physical relation and contextual properties of the world, like how a perceptual field is made up. Thus, successful registrations may be equivalent to perceptual access in a cartographic case of tracking mental states: the perceptual field is akin to a subset of a holder's map that the agent's physically-determined perceptual affordances permit her to encounter.

The fourth principle of minimal mindreading states that registration can guide action even when incorrect. In the false belief task, for example, an agent fails to witness a crucial change in an object's location, which renders her mental state incorrect. At this point, a cartographic and a minimal mindreader might diverge in their expectations about the agent's actions. The cartographic mindreader relies on relevant perceptual access to build states that guide actions; since the agent lacks the relevant perceptual access to the crucial, new state, she would also fail to form action expectations. The minimal mindreader, on the other hand, will persist with the agent's last successful registration as an actionable cause and expect her to act towards the old location, where the last successful encounter was represented. The crucial difference, here, seems to be in whether a mindreader is weighing the agent's epistemic status against an up-to-date factual reality, or if she is imbuing the agent with a state that persists in the face of a crucial change in the environment. In other words, if the strategy is reality-first, or agent-first.

Following the formulation in which incorrect registrations can still guide actions, registration might also trump change of object state or location in the justified true belief cases. Once a successful registration marks a state or a location, a change in re-baiting the same location need not cause a breakdown, as the (incorrect) registration is still useful. A cartographic commitment to perceptual access would, inversely, expect a breakdown in this situation, similar to the false belief situation. Thus, while registration representation can explain success in spontaneous behaviour false belief tasks, it seems to fail to articulate its limiting cases, i.e., why is the same mechanism not activated in 'classical' false belief tasks, or justified true belief tasks, both of which report failures in the great ape population.

Factive ToM

The factive ToM model explicitly relies on the concept of perceptual access to account for letting agents form actionable mental states. Phillips and Norby (2019) include two main ingredients in their account: tracking and separation. Tracking maintains the content of an agent's perspective, and separation sets it apart from the holder's own knowledge. The first criterion is compatible with the notion of perceptual access – one tracks the kind of information that goes into what and whether an agent can experience. The second turns out to be problematic: Phillips and Norby (2019) invoke the use of maps, at least as an instructive analogy (not unlike this chapter), in order to explain how information is being tracked. When an agent's perceptual access to the holder's map is tracked, they suggest that the holder makes a separate copy of their own map in order to accommodate the altercentric agent's knowledge and keep it separate from one's own. This separate map is not meant to be in a counterfactual relationship with the holder's egocentric map because it cannot add any new content, only delete the content that another agent is ignorant to, which avoids a counterfactual situation.

The separation element is problematic because, just by virtue of making an additional copy, one would presumably have to deploy a mapping function which identifies and copies all relevant information from the holder's map to the agent's. Then, there must be an additional layer that updates both maps independently and according to their changing relations. The first function is intractable because one cannot identify all possible information that constitutes a map, and likewise, can copy that relevant information into another. There is no limit as to which relevant distinctions and structures must be made and then copied. The second function, likewise, must select which information gets updated for each map that is being tracked, and keep the updating processes independent from each other. All the while, the multiple maps have to avoid incommensurability. It is not clear how the separation function avoids intractability or incommensurability.

The separation function of Phillips et al. (2021) is incompatible with the cartographic system, and further, its specifics seem psychologically intractable. Instead of making an independent, functionally separate map for an agent, it would seem more tenable for a cartographic mindreader to keep an egocentric, holistic map-like representation of the world and track the physical information that counts towards an agent's perceptual access to one's own map-like knowledge.

Making new maps for agents pertains to the notion of 'awareness relations' (Horschler et al., 2019; Martin & Santos, 2016) through which some authors attempted to account for nonhuman primate failure in false belief situations. According to awareness relations, a primate may track the relation between an agent, some information in the world, and the cues relevant to their mental states (e.g., gaze direction). Awareness relations are successful when the relation is established, much like a perceptual access or a registration would be. The authors dissociate their account from a 'knowledge-ignorance' account because, according to their formulation, primates cannot represent 'ignorance per se' which includes tracking a relation between an agent and information in a way that is decoupled from their egocentric reality. This may be compatible with how Phillips and Norby (2019) propose to separate an agent's factive map from the egocentric reality, however, Martin and Santos (2016) propose that primates do *not* do this, because doing so would be a case similar enough to false belief representation. Instead, the authors offer two operational commitments of awareness relations. First, awareness relations preclude a primate from holding positive expectations about agents being ignorant – specifically, ignorant per se, or in a relation decoupled from egocentric information. Second, primates should not act towards breaking existing awareness relations in deception or competitive situations. The first implication would need to contend with at least the positive evidence from primates' positive performance on false belief tasks, (which surfaced after their paper). Regardless of whether the primates represented the agent in a false belief scenario as ignorant 'per se' (equivalent to false belief), or ignorant as lacking knowledge, false beliefs would be outwith factive mindreaders' abilities.

The second prediction which states that primates should be able to prevent, but not also break an awareness relation is more intriguing. It is not exactly clear what breaking the relation would entail, but the authors find support in findings from Karg et al. (2015a) in which the apes a) failed to actively hide food from a competitor, and b) revealed less food to a knowledgeable competitor than a cooperator. However, failing to hide food is not a failure in breaking a relationship, but a prevention of a relation being established (which is compatible with awareness relation). In the latter behaviour, they could not break the experimenters' knowledge states without changing the relevant state in the world (e.g., removing the food). When a primate tracks an awareness relation, she can act on her tracking, but not also the agent's state.

Prevention of access (as in Karg et al., 2015), is compatible with (cartographic) perceptual access mechanism, as well as with awareness relation. The second implication of awareness relation – breaking the agent's awareness – is unclear. It may be that specific operational and conceptual commitments made by either Phillips and Norby (2019) or Martin and Santos (2016) on top of the cartographic perceptual access formulation are productive or necessary.

Factive theory of mind proposal which leverages the perceptual access mechanism can be subsumed under the cartographic formulation. Other more nuanced commitments about separation (J. Phillips et al., 2021), or active relation breakdown (Martin & Santos, 2016) seem to run against the basic cartographic formulation but likely fail to achieve tractability.

Mindmapping

Boyle (2019) offers a more restricted version of a cartographic ToM formulation for animals like great apes. She uses a ‘basic map’ representational format which can represent spatial properties of things in the world, but is more restricted than what I outlined above: its rudimentary functions cannot represent non-spatiotemporal properties like colour, age or weight of entities, etc. Basic altercentric maps can individuate objects that have different colours but only based on their distinct locations – a red apple is in the left corner while the green one is on the right. A subject has the relevant property information but does not extend it to other agents’ basic mental maps. The cartographic system outlined in this chapter is a less restrictive formulation than Boyle’s basic maps, in that it makes no restrictions on the type of factive information that enters a map. Basic maps, however, offer an interesting and plausible restriction, one that might turn out to be useful for consideration of cognitive efficiency or tractability.

Apart from the spatiotemporal restriction, Boyle’s model is in congruence with the present cartographic formulation on several grounds: basic maps should be committed to the isomorphism principle, the preclusion of co-reference, and bare existence quantification. The isomorphism principle here was re-articulated as the constraint of topological combinatory principle and the reliance on physical, observational information, which is in congruence with basic maps, insofar as the physical information is spatial. The constraint of co-reference pertains to cases of identity and mutually exclusive appearances so that two icons denoting mutually exclusive identities – e.g., Superman is also Clark Kent – cannot be represented in basic maps. This is the same constraint here discussed as implying incommensurability and intensional referential structure. Boyle’s constraint of bare existential quantification states that basic maps are unable to represent content as present without also including information about how – or rather where – it appears on the map. This resembles the constraint of concrete content, i.e., necessarily specifying how what is represented will look like in physical, observational terms – i.e., the constraint of concrete content. This includes the precluding of other quantifiers like SOME which do not specify observable or exact quantity, but are deliberately ambiguous as to the exact quantity estimate. As empirical verification, Boyle suggests a version of the level 2 perspective taking task, which relies on the co-reference constraint, and makes it further compatible with the present proposal. Thus, Boyle’s mindmapper model is broadly compatible with the cartographic constraints, albeit as a more restrictive version.

Cartographic theory of mind

Cartographic constraints of theory of mind are compatible with primates' successful knowledge representation, as well as failure in belief representation, but not also the belief-congruent spontaneous behaviour results. Insofar as the 'classical' and spontaneous measure tests are treated as equivalent in their theoretical implications, they will cause a conflict within any theoretical framework. There are two ways out of this. Either reconsidering how informative spontaneous tests of belief representation are (Poulin-Dubois et al., 2018), or dissociating their theoretical implications from classical tests. Tomasello (Horschler et al., 2020; 2018), e.g., offers that spontaneous false belief tasks measure representing intentional attitudes, but not propositional beliefs, difference being analogous to that between taking and confronting perspectives. Only the latter entails confronting incompatibility in an intensional referential context. Taking a perspective, on the other hand, matches the objective, factual reality, and therefore fits the extensional referential structure.

Several accounts of nonhuman primate theory of mind are broadly compatible with the dichotomous problem space of propositional and non-propositional representational formats. In some cases, models can be subsumed under others, like awareness relation converging with the perceptual access concept. Other differences, like that between minimal and factive theory of mind, which arise from minor but consequential commitments (i.e., registrations persisting while perceptual access breaking down) can stand to be queried further – the point of divergence might not be worth balkanising the problem space. They should be examined both empirically and theoretically to determine whether the differences are psychologically meaningful. Interesting constraints appear in proposals like Boyle's mindmapper model which offers an account consistent with, but restricted from the cartographic theory of mind mechanism. From a theoretical perspective, this broad coherence between different accounts offers a psychologically plausible reading of available data.

5.4.2 Object individuation

Object individuation is an instructive domain to consider the idea of different representational formats as it affords holding an object – not usually a dynamic entity – fixed while varying the way in which it is interpreted in different formats. Developmental and comparative psychology distinguish between at least three levels of the object concept, or object individuation: based on spatiotemporal, featural, or identity information. An additional category can be provisioned for sets of salient features like face-like, or agentive stimuli which seem to carry information of a particular category, but do not include some markers of identity, like varying appearances.

Spatiotemporal individuation is the lower bound of the object concept, and it involves tracking bounded, coherent wholes moving through space, but is insensitive to featural information, i.e., colour, shape, etc., In manual search paradigms, great apes, capuchins, and rhesus macaques inform their search with spatiotemporal and featural information (Flombaum et al., 2004; Kersken et al., 2020; Mendes et al., 2008, 2011; Santos et al., 2002). Rhesus monkeys also showed a looking time pattern compatible with featural individuation in an eye-tracking paradigm (Uller et al., 1997).

Much like knowledge and knowledge-like state tracking help an organism successfully navigate the social environment, so do spatiotemporal and featural object individuation help navigate the physical environment. Both spatiotemporal and featural information are concrete in content – they contain modality-specific information which describes how the objects can be observed. To 'complete' the object concept, though, we need the concept of identity: even in everyday language, we talk about objects in abstract, less observable terms. For example, objects can have deceptive appearance, like in poor visibility, they can be too small to directly observe, like atoms, or are purely theoretical like geometrical shapes or hypothesised particles. The crucial information there is not observational or physical, and therefore it is unmappable. Like in theory of mind research, a methodological challenge here is to afford an opportunity to show non-observable representation in animals through observable means only (being constrained to nonverbal manipulations and measurements). One strategy is to leverage object transformation or contrasts between their appearing and true properties to get at non-observable object states. If incompatible appearances are confronted over a concrete item, the situation imposes an incommensurate, intensional structure. If propositional representational formats are required for abstract information, and the cartographic is limited to concrete, this is good grounds for operational dissociation.

In a manual search task, Phillips and Santos (2007) showed rhesus monkeys a baiting of a piece of coconut or apple inside a box – whole fruits were visible, and a small, ambiguous white piece of food was baited. When the piece did not match the baited fruit (the piece was surreptitiously switched in some trials), the monkeys searched the box longer than when the piece and the fruit matched. While compatible with tracking whether essential and surface features match, it is possible to represent the two co-extensively so that the taste is associated alongside, and not subsumed under the fruit's identity. A more stringent test of identity individuation involves the transformation of surface properties while internal identity persists. Phillips et al. (2010) then extended the task to include the transformation of a fruit (a coconut covered with an apple shell, or vice versa), which they then baited in a box and allowed the monkeys to search in it. The monkeys' search behaviour revealed that they expected the piece they find to match the internal fruit, pre-transformation. The pattern persisted when the transformation and baiting were reversed so that it was revealed that the coconut was actually an apple (or vice-versa), eliminating the possibility that monkeys focused on the order of presentation (i.e., expecting the piece to taste like the fruit they see first).

However, the Phillips et al. (2010) transformation event may only amount to an occlusion event which need not imply an intensional or identity situation because, for the monkeys, the occlusion need not have transformed the fruit's internal identity, but only added a new surface property. Cacchione et al. (2016) attempted a stronger transformation manipulation in which experimenters took more palatable banana slices and less palatable carrot slices, and transformed one of the items in view of a sample of chimpanzees, bonobos, and orangutans. In majority of cases, the apes chose the deceptive banana (really, a carrot) more often than a virtual carrot (really, a banana). However, in a modified test with larger food items (presumably increasing motivation), the apes chose a preference for the essential banana. Interestingly, in a closely matched test where a lexigram-trained bonobo Kanzi pointed to a true (pre-transformation) property when prompted by a corresponding lexigram. Kanzi outperformed the Cacchione et al. (2016) sample. Importantly, this type of transformation task is vulnerable to a strategy akin to spatiotemporal tracking – the items do not change positions throughout the test, so the apes can focus on the location of the 'correct' item, ignoring featural transformation. In the latter, Kanzi would have to track both locations, as he did not know which type of item he should point to until after the transformation. Thus, tracking enduring essential properties in face of transformations might be within primates' abilities, but performance is not compelling, and is vulnerable to alternative explanations.

Another set of identity-involving relevant operational situations are appearance-reality tests. This was more thoroughly covered in the introduction of chapter 2 of this thesis, but suffice to say, when appearances of objects are manipulated, e.g., with magnifying lenses or coloured filters, chimpanzees show somewhat mixed performance, with high inter-individual and inter-task variability, or seem to perform well only after much exposure (Krachun, Call, et al., 2009; Krachun et al., 2016). Other tasks which show better performance need not amount to a representational situation that involves intensionality or abstract identity, as formulated here (Karg et al., 2014). In Chapter 2 here, I report on an unsuccessful replication of Krachun et al. (2016) which is not compatible with the idea that AR discrimination is a robust psychological ability in chimpanzees.

Thus, nonhuman primates' object individuation is robust in concrete, extensional cases, even in some liminal cases where occlusion may imply some low-level appearance transformation; results are mixed, however, in more conservative appearance transformation tasks, like when an entire item is painted to resemble another, or when objects' visual appearance are being manipulated. If more stringent tasks of appearance reality discrimination imply an intensional representational structure, then this may be the crucial limiting case of nonhuman primate object individuation, in line with the contrast between propositional and nonpropositional formats.

If it turns out that nonhuman primates end up not showing a hard limitation in their object individuation performance, and pass even stringent identity-based tasks, it still remains possible that egocentric object individuation does not lend itself to intensional situations with abstract content, and animals will continue to 'cheat' on AR tasks without representing the full intensional situation. Moll and Tomasello (2012) made the point of distinguishing between AR discrimination that does, or does not require confronting two incompatible aspects onto one. Younger children deploy the former strategy, effectively 'marking' deceptive or true objects as such, but not also conflicting both aspects onto one object, which only older children seem to do. This falls in line with the distinction between 'taking' and 'confronting' perspectives, whereby younger children understand which objects an actor might refer to between distorted or non-distorted objects, but do not also understand how an object appears to the actor (Moll et al., 2013; Moll & Meltzoff, 2011b). Future empirical efforts might proceed with these conceptual criteria in mind.

5.4.3 Inference

Logical operators and deductive reasoning are means of manipulating content in propositional thinking according to its relations established in formal logical operators. Nonpropositional, cartographic representations, should instead be based on observational information and inductive inference. Induction, here understood, is limited in generalising from finite observation, while deduction draws logical conclusions from premises.

Scientists regularly encounter the disconnect between deductive inference and generic statements about the world on one hand, and the need to ground them in observational terms on the other. Generic statements like 'chimpanzees have a ToM' are difficult to express in observational terms because they are 'too big' for a finite set of observations or too abstract to be confined to concrete terms. Inductive judgments about the world are probabilistic, and not conclusive or necessary. Absolute certainty or necessity is a benefit of deductive, logical inference which need not rely on (imperfect) empirical corroboration. Mathematics and logic enjoy privilege so that conclusions hold if they follow from certain premises.

Cup-game task

The way inference proceeds between premises is by use of logical operators such as IF, OR, etc., so that one can represent the disjunctive syllogism case propositionally: p OR q; NOT p; q. Translating it to an operational situation, an experimenter might present an ape with two unbaited containers, baiting both behind an opaque barrier. Then, she might reveal one container as empty, and offer them both for choice. If the ape chooses correctly, her performance is compatible with her having represented the situation in terms of the disjunctive syllogism. Of course, it is not the only plausible psychological strategy. It is possible to arrive at the same choice via, e.g., estimating the relative probability of either options and choosing the higher one. This sort of strategy is inductive as it does not require treating information as certain, but instead as options with finite probabilities. Most real-life situations where animals make choices are well suited for inductive inference, as outlined above. This again makes it hard to find a suitable, nonverbal operational situation that would inform the matter of deductive vs. inductive inference

In the basic two-cup task which may amount to the disjunctive syllogism form, the experimenter baits one of two cups behind an occluder, reveals the empty cup to the subject, and offers both up for choice. Great apes choose the correct, baited cup, above chance (Call, 2004; Engelmann et al., 2022, experiment 1). There are several alternative proposed psychological underpinnings for this pattern other than the disjunctive syllogism. The apes could have been using a heuristic of avoiding the empty, revealed cup – the 'avoid empty' strategy. Or, they might have been making serial independent guesses at the cups, i.e., making nonspecific probability judgments for each baiting location – the 'maybe A, maybe B', or 'serial guessing' strategy.

A modified three-cup design can eliminate both of these alternatives: there are three cups on the table, and two adjacent ones are occluded by a barrier. Experimenter baits one behind the barrier, so that each member of the pair has .5 chance of holding the grape. Then, she reveals one of the cups that was behind the barrier to be empty, making the other barrier cup one certain to hold the grape. Both the avoid-empty cup and the serial guessing heuristics would guide the subject to choose the correct cup in 50% of cases, but following the disjunctive syllogism, the unrevealed cup that was behind the barrier should come up as a certain outcome (Leahy & Carey, 2020; Mody & Carey, 2016). Call (2022) shows that apes perform above 50% (but well below 100%), thus showing that they are unlikely to be using either alternative heuristics. In a similar version of the 3-cup task (Engelmann et al., 2022), a food item is baited behind the barrier covering two cups, as well as under the third cup placed outside the barrier, and the choice between the three cups is offered without any reveal. The chimpanzees in their sample chose a baited cup in 51% of trials, which is compatible with the alternative strategies. In a further modification, the four-cup task, there are four cups on the table, two pairs covered by a barrier each. A food item is baited behind each barrier. Then, one empty cup is revealed, which makes the other cup of its pair certain to contain the grape, while in the other pair, it is 50% between either cup. Alternatively, under either the avoid-empty or serial guessing heuristics, one should select the target grape 1/3 of the time, as all remaining options are equally plausible. The chimpanzees chose the target in 48% of trials which is higher than the low-level heuristics, or chance, would expect, but is well below certainty levels of ~100%.

It is worth pausing to reflect on the proposed alternative strategies. The 'avoid-empty' strategy is a straightforward case of a non-inferential heuristic. There is contention to draw in its application to the three-cup task: the operational hypothesis states that the two remaining unrevealed cups should be equally likely candidates. This does not appear to be the case empirically (Call, 2022, experiment 1), but furthermore, the hypothesis and design assume that the two cups are treated equally despite one of them initially sitting behind a barrier, the other being open. Call (2022), however, shows that the apes' errors are biased towards the cup behind the barrier over the not-occluded one, which, in tandem with falsifying the hypothesis, casts doubt on its psychological plausibility. The apes may 'chunk' probabilities of outcomes based on physical constraints, so that even when they make wrong choices, they are still biased towards the barrier, presumably because they have some partial understanding that the grape must be somewhere in the area behind the barrier. Thus, the 'avoid-empty' heuristic can be discarded on grounds of lack of empirical support.

The guessing heuristic stands to be considered with respect to the cartographic nonpropositional constraints. The heuristic supposes that animals are making quantitatively non-specific guesses – akin to a 'maybe' judgment – as to the location of the food (Leahy & Carey, 2020). The 'maybe' judgment is specified in operational, but less so in conceptual terms, but it seems to resemble the SOME quantifier, which, in line with Camp (2007), is outside the scope of a cartographic system. Like SOME, a 'maybe' judgment is nonspecific in that it does not specify a finite quantitative reference. Assuming the cartographic constraints, 'maybe' can be instantiated either with some quantity estimate, inductively, but if it is nonspecific, it is likely in the deductive, propositional space. From the way Leahy and Carey used the maybe operator, it seems unlikely that it concerns specific quantities: they preclude updating probability estimates based on new information or specifying the exact quantities based on the choice space. However, it is also unlikely to be a deductive operator, given that it is launched as an alternative to deploying the OR operator and modal concepts in general. Therefore, the psychological instantiation of the serial guessing heuristic is unclear.

Another protocol suggests away from the possibility that apes are approaching quantitative situations deductively. Hanus and Call (2014) presented chimpanzees with two trays where each could contain between 1-6 opaque upturned cups, and baited between 1-4 cups on each tray, which were then shuffled under an occluder among the rest of the empty cups. The number of cups and the proportion of baited cups varied between trays and between trials so that on each trial, each tray had a different ratio of baited and empty cups. On each trial, the chimpanzees could choose a cup from one of the trays and would receive a reward if they chose a baited cup. The rates of their choices between trays correlated highly with the linear function of the difference of baited cup ratios, which means that the larger the difference in probability of finding a food item on trays, the more optimal their choices were. Meaning they were accurately tracking the relative probabilities of a particular tray holding a reward. Importantly, however, their choices were guided by relative probabilities even in cases where a tray was certain to contain a food item, i.e., the apes ignored the fact that some outcomes are 100% certain. Forgoing an 'easy' certain judgment and persisting with a quantity estimate can be taken as a characteristic limit of inductive inference, which is barred from deductive certainty judgments and committed to finite estimates.

Overall, it seems like great apes' performance on the two-cup and similar tasks cannot be explained by low-level non-inferential strategies, but their performance also fails to approach near-absolute-certainty levels, which is a mark of deductive inference. Plausibly, apes infer using quantitative probability estimates attached to baiting locations and ignore certain outcomes. Further, they may be using some spatial heuristics for constraining their choices, given, e.g., how their errors are biased towards areas behind barriers. This is broadly compatible with the cartographic constraints

Y-shaped tube task

Under the deductive psychological mechanism, the OR operator would allow for entertaining options counterfactually – e.g., 'either left or right cup must be baited' – and judge events with certainty. One other operational situation that requires counterfactual inference is the y-shaped tube task, where apes or children faced a vertical tube which forks in two ends at the bottom. In its simpler form, the experimenter drops an item into the top of the tube, which then has an equal chance of dropping out of either end at the bottom. In order to cover for both outcomes, a subject must cover both exits with their hands, or else they risk losing the reward. Great apes cover only one end, performing worse than even 2-year-olds (Redshaw & Suddendorf, 2016). They perform similarly in a modified version with two independent tubes where the experimenter dropped the items in either tube randomly (Suddendorf et al., 2017). Children, on the other hand, improve through ontogeny, with 4-year-olds reaching proficiency in the original y-shaped tube task, and somewhat poorer performance (but still better than younger children) in the experimenter-driven modified task. These data might indicate an inability to represent two possibilities simultaneously.

Other data challenges this notion. Engelmann et al. (2023) modified the y-shaped tube task to a competitive context in which chimpanzees faced either a y-shaped tube terminating over two trays, or a single vertical tube above a single tray. The platforms were both baited with food pieces, and a competitor experimenter could put a stone through the top of the tube, which tipped the platforms towards them, and the food piece would roll away from the chimpanzees. On each trial, the chimpanzees could stabilise the platforms in order to protect them from rolling away. They could press down on either one or both trays to secure the food from the experimenter stealing it. They pressed on both trays more often in the y-shaped tube condition compared with the single tube condition, which suggests awareness that the y-shaped tube allows the stone to fall on either platform and that they should both be protected.

In another task, Engelmann et al. (2021) baited one of two boxes on separate platforms connected by a rope, and the chimpanzees could pull either both ends of the rope to pull both boxes or one end of the rope to pull one box. Chimpanzees pulled on both rope ends when they could not see the box content compared to when they could (experiment 1), and when they did not witness the baiting of opaque containers compared to when they did (experiment 2). Thus, they acted on both possibilities when they did not know the location of the food item, showing some appreciation of uncertainty over alternative possibilities.

Thus, there are mixed evidence of great apes representing multiple possibilities. Plausibly, there is some underlying factor uniting the operational definitions of tasks that afford the positive and the negative results. I tentatively draw the distinction between competing and incompatible possibilities: competing possibilities have finite probability estimates and compete for informing the choice. Incompatible options cannot be apprehended 'on the same plane' and they are not considered co-extensively, as parts of the same probability space. This distinction would need to translate into operational terms.

We may tentatively conjecture that holding both hands up to tube ends is more difficult for the apes than either pulling on two ends of the rope or pressing on two trays simultaneously. One plausible difference between the two situations pertains to acting *on* an object, whereby pulling on a rope or pressing on a tray alters the state of an apparatus, but holding up two hands under tube ends does not. The former might stress the physicality of the situation, and lend itself better to the cartographic representational constraints, specifically, that of co-extensive concrete content. If one can translate multiple possibilities into action plans, then this might force the situation into extensional space where options compete, but are not incompatible. Conversely, holding up hands to the y-shaped tube may not constitute acting *on* the apparatus (i.e., causing a change in it), which fails to coerce options onto the same action plan, which leaves them incompatible.

There is some precedence for imposing concrete psychological strategies onto experimental situations. Christie et al. (2016) showed that both apes and children prefer direct over arbitrary spatial rules in a manual search task. When they saw a baiting in one of three shelves (top, middle, or bottom), and searched for an object in a separate, matching set of shelves, they were more successful when the second (hidden) baiting matched the first in a direct spatial way (top-top, middle-middle, bottom-bottom) than non-directly and 'arbitrarily' (top-middle, middle-bottom, bottom-top). Thus, relations were more obvious when they were spatially straightforward, compared to when they were arbitrary. If there is a tendency to draw relations in a non-arbitrary way, then a situation that supports this approach would lend itself to better performance. Making a task like the y-shaped tube task more actionable or embodied, i.e., allowing an ape to physically manipulate the situation, might lend itself more to a ‘concrete’ representation. Far from being a mature operational formulation, it is a plausible interpretation that might confer validity to tasks that show negative and positive performance alike.

One extenuating factor is the validity of the tube task: Lambert and Osvath (2018) showed that chimpanzees fail to cover ends of two vertical tubes, regardless of whether the outcome was uncertain as in Suddendorf et al. (2017), or certain (when that the experimenter dropped two grapes into two vertical tubes). Curiously, 4/6 subjects eventually used the method of covering both holes in later trials of the certain condition but reverted back to using just one hand. This brings the validity of the tube task data into question and further complicates the interpretation.

Regardless, great apes’ inference about events might be driven by cartographic, i.e., inductive and concrete, representational constraints. This is a cognitively ‘richer’ mechanism than the ‘low-level’ heuristics like serial guessing and ‘avoid empty cup’, but ‘leaner’ than deductive representational strategies which allow deployment of logical operators and inferring certain outcomes.

5.4.4 Communication

Since I discussed the communication problem space in earlier chapters, I will only briefly reiterate some theoretical distinctions. Gestural communication lends itself well to the study of nonhuman primate communication; experimentally, it is usually examined through pointing behaviour in object choice tasks where primates request objects or containers holding objects. According to Gomez (2009), in line with the notion of Brentainan intentionality (Crane, 1998), for a point to mean something is for it to have an imagined physical relation with a target, like an imaginary line of point between a pointing finger and an object. All mental phenomena, in this reading, are *about* something, even in the absence of an immediate physical object in the world (i.e., since the relation is imagined, it can persist in the absence of some components).

Imagined physical relations are a powerful and flexible means of establishing communicative interactions, but they may not suffice to fully explain human intentional communication in, e.g., certain Gricean readings – e.g., Wilson & Sperber (2012) – which require a reference or intention without a physical referent. Under some readings, the key ingredient of intentional communication is the intention to communicate, a second order intention directed at the recipient, and not a referent in the world. In the Brentanian reading of intentional availability, relations between agents and their referents are maintained with or without an immediately present referent. However, this alone does not amount to Gricean intention to communicate which is independent of overt reference, at least not without an upgrade to intentional nonexistence, which is a more disembodied intentional kind of relation. While the referent in intentional inexistence can be maintained in the absence of some relevant content, in reality, a nonexistent referent can be maintained counter to the real state of affairs. This might be a structural equivalent of a Gricean intention to communicate, which also does not depend on reference to extensional entities.

It is not obvious how nonexistent intentionality should manifest in a manual pointing situation, as manual points are always embedded in physical relations. Results from Bohn (2015) show that apes can point to empty containers which normally hold a preferred food item; however, this does not imply nonexistence, only inexistence, as the referent is only absent from the scene, and is still implicitly referred to. In chapter 3, I developed an operational situation which purported to disentangle inexistent from nonexistent (in Brentanian terms), or first- from second-order intention (in Gricean terms). The chimpanzee sample showed sensitivity to first, but not second-order reference, which converges with a constraint of extensional, referential structure of cartographic representation.

Summary of empirical proof of concept

The conceptual distinctions made at the domain-general level can be reduced to operational, domain-specific terms. For the theory of mind, much of the theoretical work has already been developed by other authors; the contrast between propositional beliefs and non-propositional attitudes was already prevalent in the literature, and many theories converged on similar dichotomies. In line with the theory integration approach, I weighed them against the domain's general cartographic constraints pertaining to extensional reference.

Apes' performance in object individuation tasks maps relatively well onto the propositional-cartographic dichotomy. Spatiotemporal and featural object individuation carry domain-specific, concrete information without much room for dissociating appearance and reality, or for identity structures. Nonhuman primates readily individuate on spatiotemporal and featural grounds. Inversely, in tasks which require representing identity or contrasting appearance and reality, primates show mixed results. In the domain of inference, cup choice tasks show, does not lend themselves to low-level explanations, but are also unlikely to be deductive. Performance on the y-shaped tube task, which may require the deployment of counterfactuals, yields mixed results.

For the domain of communication, I briefly reiterated the theoretical points from Chapter 3, to conserve space. Accounts of both Brentanian and Gricean intentional communication divide in forms of intention more and less reliant on overt referents in the world, the former of which converges with the concrete content and extensional referential structure of the cartographic system.

5.5 Conclusion

The prevalent debate in comparative cognition has long revolved around whether human and animal minds should be explained in the same cognitive, terms. Usually, the debates are stretched between the ends of associactivist or 'low-level' explanations and cognitivist explanations, representing typically animal and typically human cognition, respectively. If we accept that the propositional LoT-esque model was the default position of the cognitivist explanatory camp, and the associativist akin to behaviourism, then we might also consider that these options are stretched too far apart, and the problem space might benefit from introducing a third, intermediate, position. The cluster of theories that appear around the minimal, cartographic, and iconic concepts are good candidates. Wealth of nonhuman primate cognitive experiments suggest some level of cognitive ability and flexibility, but also a degree of performance limitations which stand to be accounted for in a coherent way. The theoretical sketch outlined in the present chapter makes room for such a solution. Maps and the cartographic representational system maintain the discontinuity with human cognition, all the while affording a powerful cognitivist explanatory language for nonhuman primate, and specifically chimpanzee, cognition.

Chapter 6: General discussion

6.1 Introduction

The guiding heuristic of this project has been the contrast between 'external' and 'internal' psychological representation, or between representing observable and physical content on the one hand, and the mental and internal on the other. This dualist problem space is expressed in the basic formulation of a mental representation by which an internal state can be about an external referent. A representation is, regardless, decoupled from its external referent, and can be about things that cannot be expressed in physical, observable terms. On the other hand, associativist accounts are non-representational counterparts that imply a more straightforward concurrence or extension between an external and an internal state does not have the same representational function built-in but can be described in, e.g., associativist terms. This dichotomy stretches the problem space far apart between characterising the mind in either cognitivist, representational terms or associativist, non-representational terms.

Certain cognitivist theories purport to carve a space for broadly cognitivist descriptions of the mind that exhibits certain limitations in comparison to propositional, human-typical representational abilities. This explanatory strategy seems well placed to explain nonhuman animal cognition, at least that of great apes, or chimpanzees, who demonstrate rich cognitive abilities but still exhibit characteristic limitations compared to humans. In this PhD project, I applied some operational implications of such accounts in the domain of theory of mind, and I developed a more comprehensive system of a proprietary set of domain-general cognitive constraints of great ape cognition.

In this concluding chapter, I will make less formal remarks on the implications and issues of what I reported on in previous chapters. I start with reflecting on the predicament of leveraging perceptual information in making inferences about extraperceptual cognitive processes, as I have attempted in the appearance-reality discrimination task in Chapter 2. In further experiments, I turned to a different methodological formula. Instead of conflicting different appearances of a physical referent, in chapter 3, I conflicted instrumental and specifically communicative intentions by violating chimpanzees' requests in specific ways. In Chapter 4, I adapted a communicative task purporting to test chimpanzees' ability of communicative disambiguation of a manual act. Both of the protocols in chapters 3 and 4 relied on pragmatic assumptions which may or may not have been adequately met.

Finally, in Chapter 5, I extended the two-system explanatory heuristic into a tractable domain-general sketch of a cognitive architecture. I conjectured that there are two psychologically interesting representational formats: the propositional language of thought (LoT), and the nonpropositional cartographic format. The latter might be well-equipped to explain presumed characteristic limitations of nonhuman animals, particularly great apes. The propositional format, I conjectured, is a better fit for (mature and typical) human cognition. This divide implies a phylogenetic developmental component, but I was selectively mute on the developmental component. Here, I will briefly consider what, from the theoretical side, would be necessary to include the dynamic, developmental component in the architectural sketch.

6.2 AR discrimination: leveraging and controlling for perceptual information

Nonverbal cognitive tests are forced to use perceptual information to make inferences about extraperceptual psychological processes. The validity of this methodological strategy has been discussed at least since the appearance of violation of expectation (VoE) methods which exploit looking time to (presumably) novel types of stimuli (e.g., Haith, 1998; Paulus, 2022). One contention is in the fact that, in such nonverbal paradigms, any attempt at controlling for perceptual factors orthogonal to the target conceptual dimension will undoubtedly leave some perceptual variation which co-occurs with but is theoretically orthogonal to the relevant psychological dimension. The researcher is in an inopportune position of simultaneously controlling for some, while systematically varying other, correlated, perceptual information.

In Chapter 2, I conducted an appearance-reality (AR) discrimination task where I manipulated the apparent grape size, and therefore relative quality, with magnifying and minimising lenses. I sought to control for all perceptual cues other than the apparent size of the grapes, by using identical 3D printed lens boxes, visually inspecting grapes when selecting the grapes, as well as while they were manipulated in the boxes. The magnifying and minimising lenses distorted the sizes, and not also the shape of the grapes. However, chimpanzees' performance suggests that they were tracking some perceptual cues that correlate with the manipulations. None of the chimpanzees passed both the inverse and regular trial types of mixed sessions, where, in the former, the baiting contingency was inversed and the small grape appeared even smaller, and the larger grape, bigger. It is not obvious which perceptual factor contributed to the chimpanzees' discrimination – either they tracked the retinal image associated with the truly larger grape, the relative size, or a subtle factor that correlated with the lens manipulations like light refraction or viewing angle of the lenses. This alone is instructive in that live perceptual manipulations are vulnerable to intractable perceptual factors that correlate with the crucial manipulation. This intractability, then means that the alternative hypothesis is also intractable.

One way to improve the epistemic situation, though, is to reformulate the alternative – perceptual, or associativist – hypothesis into a more constrained version. In so far as associativism is akin to behaviourism, we can constrain the hypothesis with a particular reinforcement history and have it compete with a cognitivist explanation. Alternatively, one might constrain the perceptually guided hypothesis to a known perceptual effect. In any case, constraining the alternative hypothesis should prove opportune in making the alternative hypotheses more tractable.

6.3 Pragmatic factors in communicative tasks

In Chapter 3, I pitted violations of instrumental and specifically communicative intentions of chimpanzees' food requests. The chimpanzees chose between an array of high-quality (HQ) and low-quality (LQ) food items on the table, and I responded to their requests by sharing one of the items. In cases where I shared the item they did not point to, I violated their request, and depending on the exact item I shared, the violation could pertain to either their instrumental or communicative intention or to their communicative intention alone. I found evidence of sensitivity to violations of their instrumental, but not communicative violations.

In the protocol, the chimpanzees pointed towards food items to express their preference and requests. The assumption was that, in order for their communicative acts to be meaningful, the information about their preference would have to be informative, so that the recipient would not have already known what they were conveying before they pointed. However, the chimpanzees had an opportunity to demonstrate their preference (to the same experimenter) both across trials, as well as in the preference test phase of the experiment. Informing the experimenter of their preference, then, was not necessarily pragmatically meaningful, because both would have reasons to know what the other meant already. In so far as the protocol relies on the pragmatic factors of the interaction, the interpretation of the data is made less clear.

The protocol reported in chapter 4 had chimpanzees request HQ and LQ referents by pointing at them, to which I responded by sharing their preferred item. In half of the trials, the most direct line of point was ambiguous because the imaginary line of point intersected an LQ distractor before the HQ target. The chimpanzees could modify their direct default line of point to avoid intersecting the LQ, and intersect only the HQ item. Moreover, a further factor of spatial proximity of referents exacerbated the risk of their point being misconstrued, which would have exacerbated the rates of point modifications if they were guided by attempts to disambiguate their reference for the experimenter. The chimpanzees modified their points more when their direct points would have been ambiguous, but this effect was smaller and did not reach significance when the referents were positioned close together, rather than farther apart. This is incompatible with the idea that they were communicatively disambiguating, but is compatible with the idea under which increasing the proximity of the HQ referent increased the rate of modified points. Likely, the chimpanzees were tracking the spatial position of the HQ referent, but not also how the recipient of the request would interpret their request.

Like the experiment in chapter 3, this protocol also relies on the assumption that the apes' requests were pragmatically meaningful, in that they were informing the experimenter of something she did not necessarily know. Likewise, this pragmatic aspect may not have been appropriately addressed, given that the preference for the HQ referent over the LQ referent was already 'established' in previous tests and preference trials.

It is possible to address this pragmatic factor in a protocol, e.g., by having a novel experimenter receive requests and share items in the test phase compared to the preference trials. However, such a protocol would again presume the chimpanzees' understanding of pragmatics and coordinating mental states: they would have to understand that some information about their own preference is known to one, but not the other experimenter. This would presume a high degree of awareness of pragmatics and coordinating mental state discrepancies, which is, in turn, being tested in the same protocol.

6.4 Implications for chimpanzees' cognitive abilities – plausibly cartographic?

The results from the AR discrimination task, here, suggest that the apes are likely tracking some 'low level' perceptual cue that correlates with the lens manipulations, and not considering how grapes of different sizes interact with the minimising or maximising lens boxes it was baited inside of. Such performance is broadly compatible with the idea that chimpanzees are using a cartographic representational format in AR discrimination, but it is not also fully exhaustive of the abilities it offers. The cartographic format would allow for a degree of AR discrimination, but not also such that it implies conflicting incompatible aspects under which an object appears, only recognising the nature of visual distortion and tracking their interaction with the object. However, the chimpanzees' performance does not amount to this level, and can instead be explained in terms of tracking perceptual information.

In Chapter 3, I reported on finding sensitivities in violations to the chimpanzees' instrumental, but not specifically communicative intentions embedded in their requests for food items. This is compatible with their communicative intentions being restricted to the first-order, Brentanian intentions by which they are intending towards things in the physical world, but not necessarily also towards recipients themselves. This alone is in line with the cartographic constraints. However, the protocol cannot rule out the possibility that the chimpanzees are only reacting to receiving an item they do not like, i.e., their non-communicative goal of obtaining an item is not met. It is not clear how the data populates the cartographic constraints – does it fully exhaust its representational opportunities, or are 'lower' still explanations adequate?

In chapter 4, I showed that chimpanzees failed to disambiguate a referent embedded in their communicative requests, tracking instead only where the referent (a palatable food item) is positioned in space. The placement of their points (which, varying, constituted gestural modifications), was effectively an estimate of the distance of the target referent. While finite distance estimates are compatible with the cartographic constraints, this pattern does not fully exhaust the communicative repertoire of a cartographic format. Disambiguation, in the context of this protocol, could be achieved by tracking the physicalities of imagined lines of point and the recipient's perceptual access to them.

Across the empirical chapters 2-4, the data suggests something like the lower end of a cartographic representation, by which most of the cartographic representational opportunities would not be utilised. The fit between the data and the theoretical sketch is not perfect. It would have been more opportune to spend more time defining the lower end of the representational format, so that one could be either more conclusively inclusive, or exclusive, in matching up the data with the architecture.

6.5 A comment on development

Before closing, I want to reflect on a dimension of the cognitive architecture that is missing from the sketch in Chapter 5. The picture is of stable representational formats, with the more limited, cartographic area purporting to explain nonhuman, and particularly great ape, cognition, and the propositional area reserved for humans. In a more comprehensive review, I would also have included younger children's abilities under the former. Such a sketch implies a developmental component on which I have been selectively mute. The reason being, development is not straightforward to articulate in representational, cognitivist terms, and doing so would risk leaving the basic assumptions of the representational explanatory strategy.

Some readings of the basic cognitivist formulation (roughly, the mental representation) are sceptical towards accommodating for cognitive change in the same terms (Allen & Bickhard, 2013; Rey, 2014). At risk of being overly reductive, the problem is in the lack of a mapping function that would relate one cognitive structure with another, discontinuous structure which cannot be reduced to the former. Worse yet, the problem of cognitive emergence needs a function to explain how cognitive phenomena can arise from non-cognitive conditions. The most obvious solution is to propose a degree of innate content that prefigures possible change, much like a physical organ is presumably innately predisposed to mature in the body. This gambit does away with the need to explain development, or it rejects it as a possibility. I characterised this function as the ‘induction gap’ and used it to distinguish representational from non-representational structures, as well as different representational formats from one another.

While acknowledging it, I consciously neglected resolving the inductive gap. Instead, I choose to remain with the ‘orthodox’ formulation of cognition which precludes development in a strong sense. In order to address development, then, one would have to find a different formulation of the basic explanandum of cognition. Allen and Bickhard (2013) suggest moving from a representational to an interactivist formulation, which characterises the mind as a process, rather than a fixed cognitive ability. This is reminiscent of a Piagetian approach of explaining the mind in terms of processes that stabilised formation of cognitive structures, rather, representational structures themselves, which is a more cognitivist attitude. The basic formulation of an interactive theory is a process; conceivably, such a process-based theory would struggle with explaining stable structures, as cognitivism struggles with describing change. A coordinated joint endeavour between representation- and process-based explanations, then, might paint a comprehensive, compelling picture of the mind in ways that theoretical vantage could not amount to.

6.6 Closing

Just like the chimpanzee may struggle with handling belief states, so can the researcher struggle with handling different theories of the chimpanzees’ theories of mind. I chose to constrain the problem space to a basic dichotomy, and show that the heuristic converges with the data – or at least with the disciplinary habits of cognitive psychologists. Much of what I reported here now seems, at best, as proof of concept, but maybe that is because of the scale of the problem I opened up the space for. Big and interesting problems are good goals to fall short of.

References

Allen, J. W. P., & Bickhard, M. H. (2013). Stepping off the pendulum: Why only an action-based approach can transcend the nativist–empiricist debate. *Cognitive Development*, *28*(2), 96–133. https://doi.org/10.1016/j.cogdev.2013.01.002

Apperly, I. A., & Butterfill, S. A. (2009). Do humans have two systems to track beliefs and belief-like states? *Psychol Rev*, *116*(4), 953–970. https://doi.org/10.1037/a0016923

Apperly, I. A., & Robinson, E. J. (2003). When can children handle referential opacity? Evidence for systematic variation in 5- and 6-year-old children's reasoning about beliefs and belief reports. *Journal of Experimental Child Psychology*, *85*(4), 297–311. https://doi.org/10.1016/S0022-0965(03)00099-7

Baillargeon, R., Buttelmann, D., & Southgate, V. (2018). Invited Commentary: Interpreting failed replications of early false-belief findings: Methodological and theoretical considerations. *Cognitive Development*, *46*, 112–124. https://doi.org/10.1016/j.cogdev.2018.06.001

Bohn, M., Call, J., & Tomasello, M. (2015). Communication about absent entities in great apes and human infants. *Cognition*, *145*, 63–72. https://doi.org/10.1016/j.cognition.2015.08.009

Boyle, A. (2019). Mapping the Minds of Others. *Review of Philosophy and Psychology*, *10*(4), 747–767. https://doi.org/10.1007/s13164-019-00434-z

Boysen, S. T., & Berntson, G. G. (1995). Responses to Quantity: Perceptual Versus Cognitive Mechanisms in Chimpanzees (Pan Troglodytes). *Journal of Experimental Psychology: Animal Behavior Processes,* *21*(1), 82.

Boysen, S. T., Berntson, G. G., & Mukobi, K. L. (2001). Size matters: Impact of item size and quantity on array choice by chimpanzees (Pan troglodytes). *Journal of Comparative Psychology*, *115*(1), 106–110. https://doi.org/10.1037/0735-7036.115.1.106

Bräuer, J., Call, J., & Tomasello, M. (2005). All Great Ape Species Follow Gaze to Distant Locations and Around Barriers. *Journal of Comparative Psychology*, *119*(2), 145–154. https://doi.org/10.1037/0735-7036.119.2.145

Bräuer, J., Call, J., & Tomasello, M. (2007). Chimpanzees really know what others can see in a competitive situation. *Animal Cognition*, *10*(4), 439–448. https://doi.org/10.1007/s10071-007-0088-1

Bulloch, M. J., Boysen, S. T., & Furlong, E. E. (2008). Visual attention and its relation to knowledge states in chimpanzees, Pan troglodytes. *Animal Behaviour*, *76*(4), 1147–1155. https://doi.org/10.1016/j.anbehav.2008.01.033

Buttelmann, D., Buttelmann, F., Carpenter, M., Call, J., & Tomasello, M. (2017). Great apes distinguish true from false beliefs in an interactive helping task. *PLOS ONE*, *12*(4), e0173793. https://doi.org/10.

1371/journal.pone.0173793

Buttelmann, D., Carpenter, M., & Tomasello, M. (2009). Eighteen-month-old infants show false belief understanding in an active helping paradigm. *Cognition*, *112*(2), 337–342. https://doi.org/10.1016/j.cognition.2009.05.006

Butterfill, S. A., & Apperly, I. A. (2013). How to Construct a Minimal Theory of Mind. *Mind & Language*, *28*(5), 606–637. https://doi.org/10.1111/mila.12036

Cacchione, T., Hrubesch, C., Call, J., & Rakoczy, H. (2016). Are apes essentialists? Scope and limits of psychological essentialism in great apes. *Anim Cogn*, *19*(5), 921–937. https://doi.org/10.1007/s10071-016-0991-4

Call, J. (2004). Inferences About the Location of Food in the Great Apes (Pan paniscus, Pan troglodytes, Gorilla gorilla, and Pongo pygmaeus). *Journal of Comparative Psychology*, *118*(2), 232–241. https://doi.org/10.1037/0735-7036.118.2.232

Call, J. (2022). The “avoid the empty cup” hypothesis does not explain great apes’ (Gorilla gorilla, Pan paniscus, Pan troglodytes, Pongo abelii) responses in two three-cup one-item inference by exclusion tasks. *Journal of Comparative Psychology*. https://doi.org/10.1037/com0000321

Call, J., Agnetta, B., & Tomasello, M. (2000). Cues that chimpanzees do and do not use to find hidden objects. *Anim Cogn*, *3*(1), 23–34. https://doi.org/10.1007/s100710050047

Call, J., Hare, B. A., & Tomasello, M. (1998). Chimpanzee gaze following in an object-choice task. *Anim Cogn*, *1*(2), 89–99. https://doi.org/10.1007/s100710050013

Call, J., Hare, B., Carpenter, M., & Tomasello, M. (2004). ’Unwilling’ versus ’unable’: Chimpanzees’ understanding of human intentional action. *Developmental Sci*, *7*(4), 488–498. https://doi.org/10.1111/j.1467-7687.2004.00368.x

Call, J., & Tomasello, M. (1994). Production and Comprehension of Referential Pointing by Orangutans. *Journal of Comparative Psychology*, *108*(4), 307–317.

Call, J., & Tomasello, M. (1999). A Nonverbal False Belief Task: The Performance of Children and Great Apes. *Child Development*, *70*(2), 381–395. https://doi.org/10.1111/1467-8624.00028

Call, J., & Tomasello, M. (2008). Does the chimpanzee have a theory of mind? 30 years later. *Trends in Cognitive Sciences*, *12*(5), 187–192. https://doi.org/10.1016/j.tics.2008.02.010

Camp, E. (2007). Thinking With Maps. *Philosophical Perspectives*, *21*(1), 145–182. https://doi.org/10.1111/j.1520-8583.2007.00124.x

Cartmill, E. A., & Byrne, R. W. (2007). Orangutans Modify Their Gestural Signaling According to Their Audience’s Comprehension. *Current Biology*, *17*(15), 1345–1348. https://doi.org/10.1016/j.cub.2007.06.069

Christie, S., Gentner, D., Call, J., & Haun, D. B. M. (2016). Sensitivity to Relational Similarity and Object Similarity in Apes and Children. *Current Biology*, *26*(4), 531–535. https://doi.org/10.1016/j.cub.2015.12.054

Cole, G. G., & Millett, A. C. (2019). The closing of the theory of mind: A critique of perspective-taking. *Psychonomic Bulletin & Review*, *26*(6), 1787–1802. https://doi.org/10.3758/s13423-019-01657-y

Crane, T. (1998). Intentionality as the Mark of the Mental. *Royal Institute of Philosophy Supplement*, *43*, 229–251. https://doi.org/10.1017/S1358246100004380

Dennett, D. C. (1978). Beliefs about beliefs [P&W, SR&B]. *Behavioral and Brain Sciences*, *1*(4), 568–570. https://doi.org/10.1017/S0140525X00076664

Doherty, M., & Perner, J. (1998). Metalinguistic awareness and theory of mind: Just two words for the same thing? *Cognitive Development*, *13*(3), 279–305. https://doi.org/10.1016/S0885-2014(98)90012-0

Engelmann, J. M., Haux, L. M., Völter, C., Schleihauf, H., Call, J., Rakoczy, H., & Herrmann, E. (2022). Do chimpanzees reason logically? *Child Development*, 1–15. https://doi.org/10.1111/cdev.13861

Engelmann, J. M., Völter, C. J., Goddu, M. K., Call, J., Rakoczy, H., & Herrmann, E. (2023). Chimpanzees prepare for alternative possible outcomes. *Biology Letters*, *19*(6), 20230179. https://doi.org/10.1098/rsbl.2023.0179

Engelmann, J. M., Völter, C. J., O'Madagain, C., Proft, M., Haun, D. B. M., Rakoczy, H., & Herrmann, E. (2021). Chimpanzees consider alternative possibilities. *Current Biology*, *31*(20), R1377–R1378. https://doi.org/10.1016/j.cub.2021.09.012

Fabricius, W. V., Boyer, T. W., Weimer, A. A., & Carroll, K. (2010). *True or False: Do 5-Year-Olds Understand Belief?* *46*(6), 1402–1416.

Flavell, J. H. (1986). The Development of Children's Knowledge About the Appearance-Reality Distinction. *American Psychologist*, *41*(4), 418–425.

Flavell, J. H., Flavell, E. R., & Green, F. L. (1983). Development of the appearance-reality distinction. *Cognitive Psychology*, *15*(1), 95–120. https://doi.org/10.1016/0010-0285(83)90005-1

Flombaum, J. I., Kundey, S. M., Santos, L. R., & Scholl, B. J. (2004). Dynamic Object Individuation in Rhesus Macaques: A Study of the Tunnel Effect. *Psychological Science*, *15*(12), 795–800. https://doi.org/10.1111/j.0956-7976.2004.00758.x

Fodor, J. A. (1975). *The Language of Thought*. Harvard University Press.

Fodor, J. A. (2008). *Lot 2: The Language of Thought Revisited* (J. A. Fodor, Ed.). Oxford University Press.

Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, *28*(1), 3–71. https://doi.org/10.1016/0010-0277(88)90031-5

Gigerenzer, G. (2017). A theory integration program. *Decision*, *4*(3), 133–145. https://doi.org/10.1037/

dec0000082

Gómez, J.-C. (2008). The Evolution Of Pretence: From Intentional Availability To Intentional Non-Existence. *Mind & Language*, *23*(5), 586–606. https://doi.org/10.1111/j.1468-0017.2008.00353.x

Gómez, J.-C. (2009). Embodying meaning: Insights from primates, autism, and Brentano. *Neural Networks*, *22*(2), 190–196. https://doi.org/10.1016/j.neunet.2009.01.010

Gonseth, C., Kawakami, F., Ichino, E., & Tomonaga, M. (2017). The higher the farther: Distance-specific referential gestures in chimpanzees (Pan troglodytes). *Biology Letters*, *13*(11), 20170398. https://doi.org/10.1098/rsbl.2017.0398

Gopnik, A., & Astington, J. W. (1988). Children's Understanding of Representational Change and Its Relation to the Understanding of False Belief and the Appearance-Reality Distinction. *Child Development*, *59*(1), 26–37.

Gopnik, A., & Wellman, H. M. (1992). Why the Child's Theory of Mind Really Is a Theory. *Mind & Language*, *7*(1-2), 145–171. https://doi.org/10.1111/j.1468-0017.1992.tb00202.x

Grosse, G., Behne, T., Carpenter, M., & Tomasello, M. (2010). Infants communicate in order to be understood. *Developmental Psychology*, *46*(6), 1710–1722. https://doi.org/10.1037/a0020727

Haith, M. M. (1998). Who put the cog in infant cognition? Is rich interpretation too costly? *Infant Behavior and Development*, *21*(2), 167–179. https://doi.org/10.1016/S0163-6383(98)90001-7

Hanus, D., & Call, J. (2014). When maths trumps logic: Probabilistic judgements in chimpanzees. *Biol. Lett.*, *10*(12), 20140892. https://doi.org/10.1098/rsbl.2014.0892

Hare, B., Call, J., Agnetta, B., & Tomasello, M. (2000). Chimpanzees know what conspecifics do and do not see. *Animal Behaviour*, *59*(4), 771–785. https://doi.org/10.1006/anbe.1999.1377

Hare, B., Call, J., & Tomasello, M. (2001). Do chimpanzees know what conspecifics know? *Animal Behaviour*, *61*(1), 139–151. https://doi.org/10.1006/anbe.2000.1518

Hayashi, T., Akikawa, R., Kawasaki, K., Egawa, J., Minamimoto, T., Kobayashi, K., Kato, S., Hori, Y., Nagai, Y., Iijima, A., Someya, T., & Hasegawa, I. (2020). Macaques Exhibit Implicit Gaze Bias Anticipating Others' False-Belief-Driven Actions via Medial Prefrontal Cortex. *Cell Reports*, *30*(13), 4433–4444.e5. https://doi.org/10.1016/j.celrep.2020.03.013

Heyes, C. (1998). Theory of mind in nonhuman primates. *Behavioral and Brain Sciences*, *21*(1), 101–114. https://doi.org/10.1017/S0140525X98000703

Heyes, C. (2017). Apes Submentalise. *Trends in Cognitive Sciences*, *21*(1), 1–2. https://doi.org/10.1016/j.tics.2016.11.006

Horschler, D. J., MacLean, E. L., & Santos, L. R. (2020). Do Non-Human Primates Really Represent Others' Beliefs? *Trends in Cognitive Sciences*, *24*(8), 594–605. https://doi.org/10.1016/j.tics.2020.05.009

Horschler, D. J., Santos, L. R., & MacLean, E. L. (2019). Do non-human primates really represent others' ignorance? A test of the awareness relations hypothesis. *Cognition*, *190*, 72–80. https://doi.org/10.1016/j.cognition.2019.04.012

Horschler, D. J., Santos, L. R., & MacLean, E. L. (2021). How do non-human primates represent others' awareness of where objects are hidden? *Cognition*, *212*, 104658. https://doi.org/10.1016/j.cognition.2021.104658

Hostetter, A. B., College, B., & Hopkins, W. D. (2001). Differential Use of Vocal and Gestural Communication by Chimpanzees (Pan troglodytes) in Response to the Attentional Status of a Human (Homo sapiens). *Journal of Comparative Psychology*, *114*(4), 337–343.

Hostetter, A. B., Russell, J. L., Freeman, H., & Hopkins, W. D. (2006). Now you see me, now you don't: Evidence that chimpanzees understand the role of the eyes in attention. *Animal Cognition*, *10*(1), 55–62. https://doi.org/10.1007/s10071-006-0031-x

Itakura, S., Agnetta, B., Hare, B., & Tomasello, M. (1999). Chimpanzee Use of Human and Conspecific Social Cues to Locate Hidden Food. *Developmental Science*, *2*(4), 448–456. https://doi.org/10.1111/1467-7687.00089

Itakura, S., & Tanaka, M. (1998). Use of experimenter-given cues during object-choice tasks by chimpanzees (Pan troglodytes), an orangutan (Pongo pygmaeus), and human infants (Homo sapiens). *J Comp Psychol*, *112*(2), 119–126. https://doi.org/10.1037/0735-7036.112.2.119

Kaminski, J., Call, J., & Tomasello, M. (2004). Body orientation and face orientation: Two factors controlling apes? Begging behavior from humans. *Animal Cognition*, *7*(4), 216–223. https://doi.org/10.1007/s10071-004-0214-2

Kaminski, J., Call, J., & Tomasello, M. (2008). Chimpanzees know what others know, but not what they believe. *Cognition*, *109*(2), 224–234. https://doi.org/10.1016/j.cognition.2008.08.010

Kampis, D., Karman, P., Csibra, G., Southgate, V., & Hernik, M. (2020). *A two-lab direct replication attempt of Southgate, Senju, & Csibra (2007).* PsyArXiv. https://doi.org/10.31234/osf.io/gzy26

Kano, F., Krupenye, C., Hirata, S., Tomonaga, M., & Call, J. (2019). Great apes use self-experience to anticipate an agent's action in a false-belief test. *Proc Natl Acad Sci USA*, 201910095. https://doi.org/10.1073/pnas.1910095116

Kano, F., Moore, R., Krupenye, C., Hirata, S., Tomonaga, M., & Call, J. (2018). Human ostensive signals do not enhance gaze following in chimpanzees, but do enhance object-oriented attention. *Anim Cogn*, *21*(5), 715–728. https://doi.org/10.1007/s10071-018-1205-z

Karg, K., Schmelz, M., Call, J., & Tomasello, M. (2014). All great ape species (Gorilla gorilla, Pan paniscus, Pan troglodytes, Pongo abelii) and two-and-a-half-year-old children (Homo sapiens) discriminate appearance from reality. *Journal of Comparative Psychology*, *128*(4), 431–439. https://doi.org/10.1037/a0037385

Karg, K., Schmelz, M., Call, J., & Tomasello, M. (2015a). Chimpanzees strategically manipulate what others can see. *Anim Cogn*, *18*(5), 1069–1076. https://doi.org/10.1007/s10071-015-0875-z

Karg, K., Schmelz, M., Call, J., & Tomasello, M. (2015b). The goggles experiment: Can chimpanzees use self-experience to infer what a competitor can see? *Animal Behaviour*, *105*, 211–221. https://doi.org/10.1016/j.anbehav.2015.04.028

Karg, K., Schmelz, M., Call, J., & Tomasello, M. (2016). Differing views: Can chimpanzees do Level 2 perspective-taking? *Anim Cogn*, *19*(3), 555–564. https://doi.org/10.1007/s10071-016-0956-7

Karin-D'Arcy, R. M., & Povinelli, D. J. (2002). Do Chimpanzees Know What Each Other See? A Closer Look. *International Journal of Comparative Psychology*, *15*(1). https://doi.org/10.46867/C4XW2M

Kersken, V., Zhang, D., Gomez, J.-C., Seed, A. M., & Ball, D. N. (2020). *Capuchin monkeys individuate objects based on spatio-temporal and property/kind information : Evidence from looking and reaching measures*. https://doi.org/10.26451/abc.07.03.06.2020

Kovács, Á. M., Teglas, E., & Endress, A. D. (2010). The Social Sense: Susceptibility to Others' Beliefs in Human Infants and Adults. *Science*, *330*(6012), 1830–1834. https://doi.org/10.1126/science.1190792

Krachun, C., Call, J., & Tomasello, M. (2009). Can chimpanzees (Pan troglodytes) discriminate appearance from reality? *Cognition*, *112*(3), 435–450. https://doi.org/10.1016/j.cognition.2009.06.012

Krachun, C., Carpenter, C. M., Call, J., & Tomasello, M. (2010). A New Change-of-Contents False Belief Test: Children and Chimpanzees Compared. *International Journal of Comparative Psychology*, *23*(2).

Krachun, C., Carpenter, M., Call, J., & Tomasello, M. (2009). A competitive nonverbal false belief task for children and apes. *Developmental Science*, *12*(4), 521–535. https://doi.org/10.1111/j.1467-7687.2008.00793.x

Krachun, C., Lurz, R., Russell, J. L., & Hopkins, W. D. (2016). Smoke and mirrors: Testing the scope of chimpanzees' appearance–reality understanding. *Cognition*, *150*, 53–67. https://doi.org/10.1016/j.cognition.2016.01.012

Krupenye, C., Kano, F., Hirata, S., Call, J., & Tomasello, M. (2016). Great apes anticipate that other individuals will act according to false beliefs. *Science*, *354*(6308), 110–114. https://doi.org/10.1126/science.aaf8110

Krupenye, C., Kano, F., Hirata, S., Call, J., & Tomasello, M. (2017). A test of the submentalizing hypothesis: Apes' performance in a false belief task inanimate control. *Communicative & Integrative Biology*, *10*(4), e1343771. https://doi.org/10.1080/19420889.2017.1343771

Kulke, L., & Rakoczy, H. (2018). Implicit Theory of Mind – An overview of current replications and nonreplications. *Data in Brief*, *16*, 101–104. https://doi.org/10.1016/j.dib.2017.11.016

Lambert, M. L., & Osvath, M. (2018). Comparing chimpanzees' preparatory responses to known and unknown future outcomes. *Biology Letters*, *14*(9), 20180499. https://doi.org/10.1098/rsbl.2018.0499

Langer, S. K. (1957). *Philosophy in a New Key: A Study in the Symbolism of Reason, Rite, and Art, Third Edition*. Harvard University Press.

Leahy, B. P., & Carey, S. E. (2020). The Acquisition of Modal Concepts. *Trends in Cognitive Sciences*, *24*(1), 65–78. https://doi.org/10.1016/j.tics.2019.11.004

Leavens, D. A., Russell, J. L., & Hopkins, W. D. (2005). Intentionality as Measured in the Persistence and Elaboration of Communication by Chimpanzees (Pan troglodytes). *Child Development*, *76*(1), 291–306. https://doi.org/10.1111/j.1467-8624.2005.00845.x

Liebal, K., Call, J., Tomasello, M., & Pika, S. (2004). To move or not to move: How apes adjust to the attentional state of others. *Interaction Studies*, *5*(2), 199–219. https://doi.org/10.1075/is.5.2.03lie

Lurz, R. (2009). If chimpanzees are mindreaders, could behavioral science tell? Toward a solution of the logical problem. *Philosophical Psychology*, *22*(3), 305–328. https://doi.org/10.1080/09515080902970673

Lurz, R., Krachun, C., Hopkins, W. D., & Taglialatela, J. (2022). Evidence of psychological essentialism in a symbol-trained bonobo (Pan paniscus). *International Journal of Primatology*, 1–18. https://doi.org/10.1007/s10764-021-00266-8

Lurz, R., Krachun, C., Mahovetz, L., Wilson, M. J. G., & Hopkins, W. (2018). Chimpanzees gesture to humans in mirrors: Using reflection to dissociate seeing from line of gaze. *Animal Behaviour*, *135*, 239–249. https://doi.org/10.1016/j.anbehav.2017.11.014

MacLean, E. L., & Hare, B. (2012). Bonobos and chimpanzees infer the target of another's attention. *Animal Behaviour*, *83*(2), 345–353. https://doi.org/10.1016/j.anbehav.2011.10.026

Marticorena, D. C. W., Ruiz, A. M., Mukerji, C., Goddu, A., & Santos, L. R. (2011). Monkeys represent others' knowledge but not their beliefs: Monkeys represent knowledge but not beliefs. *Developmental Science*, *14*(6), 1406–1416. https://doi.org/10.1111/j.1467-7687.2011.01085.x

Martin, A., & Santos, L. R. (2014). The origins of belief representation: Monkeys fail to automatically represent others' beliefs. *Cognition*, *130*(3), 300–308. https://doi.org/10.1016/j.cognition.2013.11.016

Martin, A., & Santos, L. R. (2016). What Cognitive Representations Support Primate Theory of Mind? *Trends in Cognitive Sciences*, *20*(5), 375–382. https://doi.org/10.1016/j.tics.2016.03.005

Mendes, N., Rakoczy, H., & Call, J. (2008). Ape metaphysics: Object individuation without language. *Cognition*, *106*(2), 730–749. https://doi.org/10.1016/j.cognition.2007.04.007

Mendes, N., Rakoczy, H., & Call, J. (2011). Primates do not spontaneously use shape properties for object individuation: A competence or a performance problem? *Anim Cogn*, *14*(3), 407–414. https://doi.org/

10.1007/s10071-010-0375-0

Mischel, W. (2008). The Toothbrush Problem. *APS Observer*, *21*.

Mody, S., & Carey, S. (2016). The emergence of reasoning by the disjunctive syllogism in early childhood. *Cognition*, *154*, 40–48. https://doi.org/10.1016/j.cognition.2016.05.012

Moll, H., & Meltzoff, A. N. (2011b). How Does It Look? Level 2 Perspective-Taking at 36 Months of Age: Level 2 Perspective-Taking. *Child Development*, *82*(2), 661–673. https://doi.org/10.1111/j.1467-8624.2010.01571.x

Moll, H., & Meltzoff, A. N. (2011a). How Does It Look? Level 2 Perspective-Taking at 36 Months of Age: Level 2 Perspective-Taking. *Child Development*, *82*(2), 661–673. https://doi.org/10.1111/j.1467-8624.2010.01571.x

Moll, H., Meltzoff, A. N., Merzsch, K., & Tomasello, M. (2013). Taking versus confronting visual perspectives in preschool children. *Developmental Psychology*, *49*(4), 646–654. https://doi.org/10.1037/a0028633

Moll, H., & Tomasello, M. (2012). Three-year-olds understand appearance and reality—just not about the same object at the same time. *Developmental Psychology*, *48*(4), 1124–1132. https://doi.org/10.1037/a0025915

Moore, C., Pure, K., & Furrow, D. (1990). Children's Understanding of the Modal Expression of Speaker Certainty and Uncertainty and Its Relation to the Development of a Representational Theory of Mind. *Child Development*, *61*(3), 722. https://doi.org/10.2307/1130957

Mulcahy, N. J., & Call, J. (2009). The performance of bonobos (Pan paniscus), chimpanzees (Pan troglodytes), and orangutans (Pongo pygmaeus) in two versions of an object-choice task. *Journal of Comparative Psychology*, *123*(3), 304–309. https://doi.org/10.1037/a0016222

Nagel, J. (2017). Factive and nonfactive mental state attribution. *Mind Lang*, *32*(5), 525–544. https://doi.org/10.1111/mila.12157

Okamoto, S., Tanaka, M., & Tomonaga, M. (2004). Looking back: The "representational mechanism" of joint attention in an infant chimpanzee (Pan troglodytes). *Japanese Psychological Research*, *46*(3), 236–245. https://doi.org/10.1111/j.1468-5584.2004.00255.x

Okamoto-Barth, S., Call, J., & Tomasello, M. (2007). Great Apes' Understanding of Other Individuals' Line of Sight. *Psychol Sci*, *18*(5), 462–468. https://doi.org/10.1111/j.1467-9280.2007.01922.x

Oktay-Gür, N., & Rakoczy, H. (2017). Children's difficulty with true belief tasks: Competence deficit or performance problem? *Cognition*, *166*, 28–41. https://doi.org/10.1016/j.cognition.2017.05.002

Onishi, K. H., & Baillargeon, R. (2005). Do 15-Month-Old Infants Understand False Beliefs? *Science*, *308*(5719), 255–258. https://doi.org/10.1126/science.1107621

Paulus, M. (2022). Should infant psychology rely on the violation-of-expectation method? Not anymore.

Infant and Child Development, *31*(1), e2306. https://doi.org/10.1002/icd.2306

Penn, D. C., Holyoak, K. J., & Povinelli, D. J. (2008). Darwin's mistake: Explaining the discontinuity between human and nonhuman minds. *BEHAVIORAL AND BRAIN SCIENCES*, 70.

Perner, J. (1991). *Understanding the Representational Mind*. MIT Press.

Phillips, J., Buckwalter, W., Cushman, F., Friedman, O., Martin, A., Turri, J., Santos, L., & Knobe, J. (2021). Knowledge before belief. *Behavioral and Brain Sciences*, *44*, e140. https://doi.org/10.1017/S0140525X20000618

Phillips, J., & Norby, A. (2019). Factive theory of mind. *Mind & Language*, *36*(1), 3–26. https://doi.org/10.1111/mila.12267

Phillips, W., & Santos, L. R. (2007). Evidence for kind representations in the absence of language: Experiments with rhesus monkeys (Macaca mulatta). *Cognition*, *102*(3), 455–463. https://doi.org/10.1016/j.cognition.2006.01.009

Phillips, W., Shankar, M., & Santos, L. R. (2010). Essentialism in the absence of language? Evidence from rhesus monkeys (Macaca mulatta). *Developmental Science*, *13*(4), F1–F7. https://doi.org/10.1111/j.1467-7687.2010.00982.x

Poss, S. R., Kuhar, C., Stoinski, T. S., & Hopkins, W. D. (2006). Differential use of attentional and visual communicative signaling by orangutans (Pongo pygmaeus) and gorillas (Gorilla gorilla) in response to the attentional status of a human. *American Journal of Primatology*, *68*(10), 978–992. https://doi.org/10.1002/ajp.20304

Poulin-Dubois, D., Rakoczy, H., Burnside, K., Crivello, C., Dörrenberg, S., Edwards, K., Krist, H., Kulke, L., Liszkowski, U., Low, J., Perner, J., Powell, L., Priewasser, B., Rafetseder, E., & Ruffman, T. (2018). Do infants understand false beliefs? We don't know yet – A commentary on Baillargeon, Buttelmann and Southgate's commentary. *Cognitive Development*, *48*, 302–315. https://doi.org/10.1016/j.cogdev.2018.09.005

Povinelli, D. J., & Eddy, T. J. (1996). Chimpanzees: Joint Visual Attention. *Psychol Sci*, *7*(3), 129–135. https://doi.org/10.1111/j.1467-9280.1996.tb00345.x

Povinelli, D. J., & Vonk, J. (2003). Chimpanzee minds: Suspiciously human? *Trends in Cognitive Sciences*, *7*(4), 157–160. https://doi.org/10.1016/S1364-6613(03)00053-6

Premack, D., & Woodruff, G. (1978). Does the chimpanzee have a theory of mind? *Behav Brain Sci*, *1*(4), 515–526. https://doi.org/10.1017/S0140525X00076512

Quilty-Dunn, J., Porot, N., & Mandelbaum, E. (2022). The Best Game in Town: The Re-Emergence of the Language of Thought Hypothesis Across the Cognitive Sciences. *Behavioral and Brain Sciences*, 1–55. https://doi.org/10.1017/S0140525X22002849

Quine, W. V. O. Q. (1951). Two Dogmas of Empiricism. *The Philosophical Review*, *60*(60), 20–43.

Rakoczy, H. (2012). Do infants have a theory of mind? *British Journal of Developmental Psychology*, *30*(1), 59–74. https://doi.org/10.1111/j.2044-835X.2011.02061.x

Rakoczy, H., Bergfeld, D., Schwarz, I., & Fizke, E. (2015). Explicit Theory of Mind Is Even More Unified Than Previously Assumed: Belief Ascription and Understanding Aspectuality Emerge Together in Development. *Child Development*, *86*(2), 486–502. https://doi.org/10.1111/cdev.12311

Redshaw, J., & Suddendorf, T. (2016). Children's and Apes' Preparatory Responses to Two Mutually Exclusive Possibilities. *Current Biology*, *26*(13), 1758–1762. https://doi.org/10.1016/j.cub.2016.04.062

Rey, G. (2014). Innate and Learned: Carey, Mad Dog Nativism, and the Poverty of Stimuli and Analogies (Yet Again). *Mind & Language*, *29*(2), 109–132. https://doi.org/10.1111/mila.12044

Roberts, A. I., Vick, S.-J., Roberts, S. G. B., & Menzel, C. R. (2014). Chimpanzees modify intentional gestures to coordinate a search for hidden food. *Nature Communications*, *5*, 3088. https://doi.org/10.1038/ncomms4088

Santos, L. R., Sulkowski, G. M., Spaepen, G. M., & Hauser, M. D. (2002). Object individuation using property/kind information in rhesus macaques (Macaca mulatta). *Cognition*, *83*(3), 241–264. https://doi.org/10.1016/S0010-0277(02)00006-9

Schmelz, M., Call, J., & Tomasello, M. (2011). Chimpanzees know that others make inferences. *Proceedings of the National Academy of Sciences*, *108*(7), 3077–3079. https://doi.org/10.1073/pnas.1000469108

Senju, A., & Csibra, G. (2008). Gaze Following in Human Infants Depends on Communicative Signals. *Current Biology*, *18*(9), 668–671. https://doi.org/10.1016/j.cub.2008.03.059

Southgate, V., Senju, A., & Csibra, G. (2007). Action Anticipation Through Attribution of False Belief by 2-Year-Olds. *Psychol Sci*, *18*(7), 587–592. https://doi.org/10.1111/j.1467-9280.2007.01944.x

Suddendorf, T., Crimston, J., & Redshaw, J. (2017). Preparatory responses to socially determined, mutually exclusive possibilities in chimpanzees and children. *Biology Letters*, *13*(6), 20170170. https://doi.org/10.1098/rsbl.2017.0170

Tauzin, T., Bohn, M., Gergely, G., & Call, J. (2020). Context-sensitive adjustment of pointing in great apes. *Sci Rep*, *10*(1), 1048. https://doi.org/10.1038/s41598-019-56183-7

Tennie, C., Call, J., & Tomasello, M. (2006). Push or Pull: Imitation vs. Emulation in Great Apes and Human Children. *Ethology*, *112*(12), 1159–1169. https://doi.org/10.1111/j.1439-0310.2006.01269.x

Tomasello, M. (2018). How children come to understand false beliefs: A shared intentionality account. *Proceedings of the National Academy of Sciences*, *115*(34), 8491–8498. https://doi.org/10.1073/pnas.1804761115

Tomasello, M., Call, J., Nagell, K., Olguin, R., & Carpenter, M. (1994). The learning and use of gestural signals by young chimpanzees: A trans-generational study. *Primates*, *35*(2), 137–154. https://doi.org/10.1007/BF02382050

Tomasello, M., Hare, B., & Agnetta, B. (1999). Chimpanzees, Pan troglodytes, follow gaze direction geometrically. *Animal Behaviour*, *58*(4), 769–777. https://doi.org/10.1006/anbe.1999.1192

Tomasello, M., Hare, B., & Fogleman, T. (2001). The ontogeny of gaze following in chimpanzees, Pan troglodytes, and rhesus macaques, Macaca mulatta. *Animal Behaviour*, *61*(2), 335–343. https://doi.org/10.1006/anbe.2000.1598

Townsend, S. W., Koski, S. E., Byrne, R. W., Slocombe, K. E., Bickel, B., Boeckle, M., Braga Goncalves, I., Burkart, J. M., Flower, T., Gaunet, F., Glock, H. J., Gruber, T., Jansen, D. A. W. A. M., Liebal, K., Linke, A., Miklósi, Á., Moore, R., van Schaik, C. P., Stoll, S., … Manser, M. B. (2017). Exorcising Grice's ghost: An empirical approach to studying intentional communication in animals. *Biological Reviews*, *92*(3), 1427–1433. https://doi.org/10.1111/brv.12289

Uller, C., Carey, S., Hauser, M, & Xu, F. (1997). Is language needed for constructing sortal concepts? A study with nonhuman primates. *Boston University Conference on Language Development Proceeding*, *21*, 665–677.

Wason, P. C. (1968). Reasoning about a Rule. *Quarterly Journal of Experimental Psychology*, *20*(3), 273–281. https://doi.org/10.1080/14640746808400161

Wellman, H. M., Cross, D., & Watson, J. (2001). Meta-Analysis of Theory-of-Mind Development: The Truth about False Belief. *Child Development*, *72*(3), 655–684. https://doi.org/10.1111/1467-8624.00304

Wiggins, D. (1997). Sortal Concepts: A Reply To Xu. *Mind & Language*, *12*(3-4), 413–421. https://doi.org/10.1111/j.1468-0017.1997.tb00081.x

Wilson, D., & Sperber, D. (2012). *Meaning and Relevance*. Cambridge University Press. https://doi.org/10.1017/CBO9781139028370

Xu, F., & Carey, S. (1996). Infants' Metaphysics: The Case of Numerical Identity. *Cognitive Psychology*, *30*(2), 111–153. https://doi.org/10.1006/cogp.1996.0005

Appendix A: Chapter 2

Figure A2.1. A reference plan used for building the lens boxes.

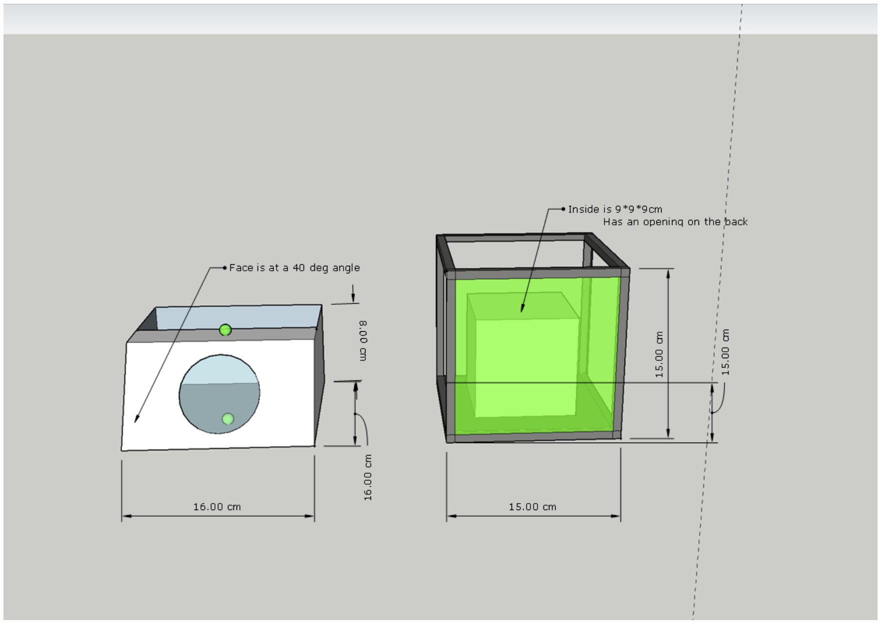

Note. The illustration also includes plans for testing apparatus for a planned colour filter test, which was not pursued.

Figure A2.2. Example of a testing sheet for the Appearance-reality testing sessions.

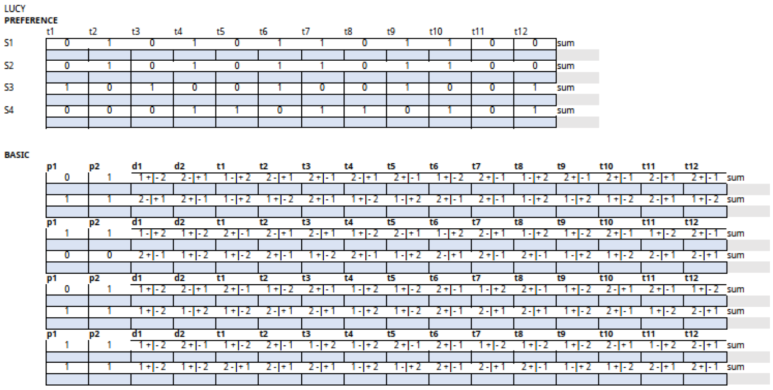

LUCY

PREFERENCE

	t1	t2	t3	t4	t5	t6	t7	t8	t9	t10	t11	t12	
S1	0	1	0	1	0	1	1	0	1	1	0	0	sum
S2	0	1	0	1	0	1	1	0	1	1	0	0	sum
S3	1	0	1	0	0	1	0	0	1	0	0	1	sum
S4	0	0	0	1	1	0	1	1	0	1	0	1	sum

BASIC

p1	**p2**	**d1**	**d2**	**t1**	**t2**	**t3**	**t4**	**t5**	**t6**	**t7**	**t8**	**t9**	**t10**	**t11**	**t12**	
0	1	1+|-2	2-|+1	1-|+2	2-|+1	2+|-1	2-|+1	2+|-1	1+|-2	2+|-1	1-|+2	2+|-1	2+|-1	2-|+1	2+|-1	sum
1	1	2-|+1	2+|-1	1-|+2	1+|-2	2+|-1	1+|-2	1-|+2	2+|-1	2+|-1	1-|+2	1-|+2	1+|-2	2-|+1	1+|-2	sum

p1	**p2**	**d1**	**d2**	**t1**	**t2**	**t3**	**t4**	**t5**	**t6**	**t7**	**t8**	**t9**	**t10**	**t11**	**t12**	
1	1	1-|+2	1+|-2	2+|-1	2-|+1	2-|+1	1+|-2	2-|+1	1-|+2	2-|+1	1-|+2	1+|-2	2+|-1	1+|-2	2+|-1	sum
0	0	2+|-1	1+|-2	1+|-2	2+|-1	1+|-2	2+|-1	1-|+2	2-|+1	2+|-1	2+|-1	1-|+2	1+|-2	2-|+1	2-|+1	sum

p1	**p2**	**d1**	**d2**	**t1**	**t2**	**t3**	**t4**	**t5**	**t6**	**t7**	**t8**	**t9**	**t10**	**t11**	**t12**	
0	1	1+|-2	2-|+1	2+|-1	1+|-2	2+|-1	1-|+2	1+|-2	2+|-1	1-|+2	2+|-1	1+|-2	2-|+1	2+|-1	1+|-2	sum
1	1	1+|-2	1-|+2	1+|-2	2-|+1	2-|+1	1-|+2	1+|-2	2+|-1	2-|+1	2-|+1	1+|-2	2+|-1	1+|-2	2-|+1	sum

p1	**p2**	**d1**	**d2**	**t1**	**t2**	**t3**	**t4**	**t5**	**t6**	**t7**	**t8**	**t9**	**t10**	**t11**	**t12**	
1	1	1+|-2	2+|-1	1+|-2	2+|-1	1-|+2	1-|+2	2+|-1	2+|-1	1+|-2	1+|-2	1+|-2	2-|+1	1-|+2	2-|+1	sum
1	1	1+|-2	1+|-2	2-|+1	2-|+1	2-|+1	1-|+2	1-|+2	2+|-1	2-|+1	2+|-1	1-|+2	1+|-2	2-|+1	2+|-1	sum

---

Note. For the seen tracking, unseen tracking, and mixed test sessions, I penciled in additional dimensions of switch/replace and standard/regular with the help of an independent randomly generated list.

Appendix B: Chapter 3

Experiment 1

Additional exploratory analysis: I re-coded the outcome variable into an alternative outcome variable which excluded pointing behaviour from the set of recorded behaviour. It is possible that the pointing gestures, which were included in the behavioural rate, served not as protest or frustration, but as a means of obtaining the remaining food items from the table – the chimpanzees might have been requesting some of the remaining food on the table after delivery. Without including pointing behaviour, the recorded behaviour is only non-instrumental.

Table B3.1 Descriptive data with respect to the outcome variable with or without pointing behaviour for Experiment 1

Condition	% of trias with any behaviour	% of trials with a non-point behaviour
Baseline	30.3%	15.96%
Identity Mistake	36.59%	21.95%
Honest Mistake	51.16%	30.23%
Offence	59.52%	35.71%

Figure B3.1 Proportion of trials with a behaviour, coded as either with or without a pointing behaviour, Ex-

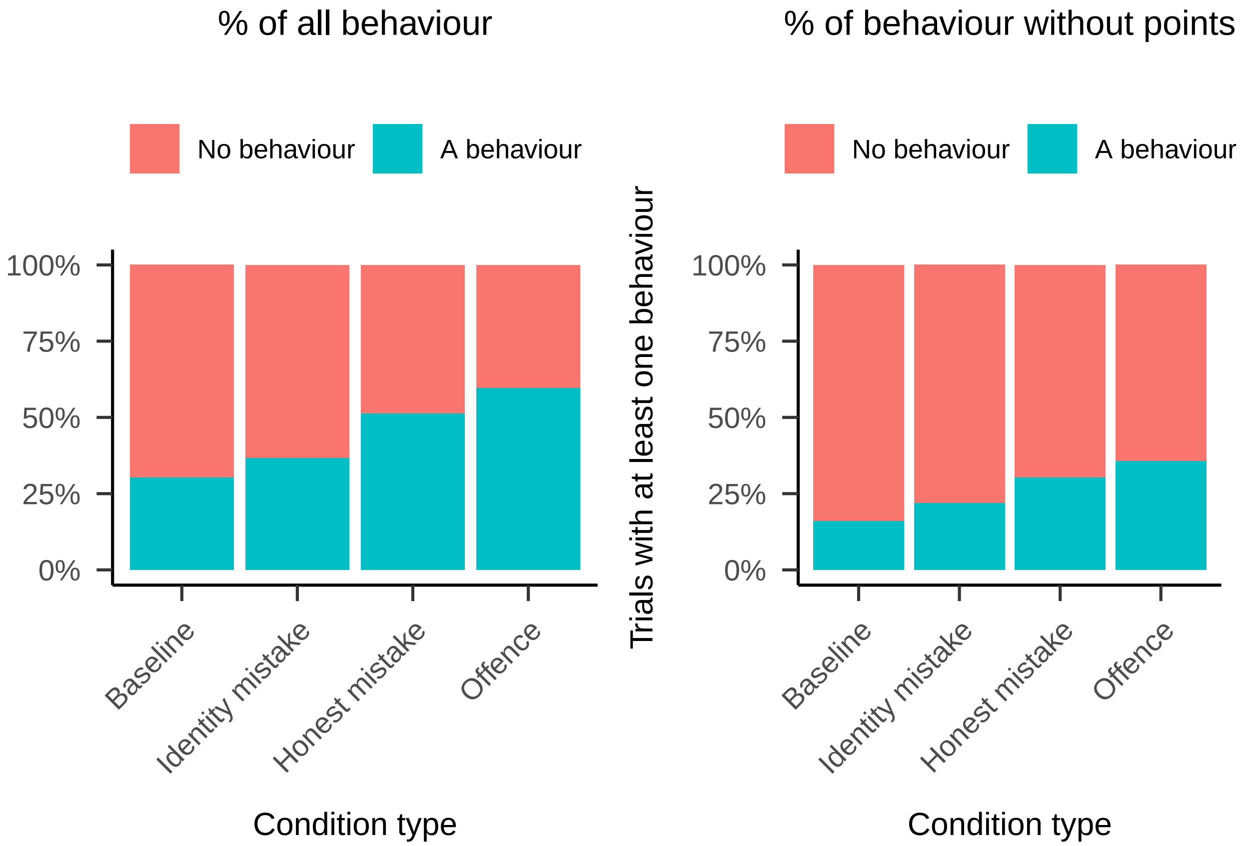

periment 1.

Experiment 2

Just like in experiment 1, I re-coded the outcome variable into 1) just pointing behaviour, and 2) all behaviour except pointing, and compared the two informally. The proportion of trials with non-pointing behaviour is summarised alongside the proportion of trials with only pointing behaviour in Table B3.2 and Figure B3.2.

Table B3.2 Descriptive data with respect to the outcome variable with or without pointing behaviour for Experiment 2

Condition	% of trias with any behaviour	% of trials with a non-point behaviour
Same Regular	23.3%	6.8%
Different Regular	3.88%	1.94%
Same Inverse	37.25%	5.88%
Different Inverse	32.69%	17.31%

Figure B3.2 Proportion of trials with a behaviour, coded as either with or without a pointing behaviour, Ex-

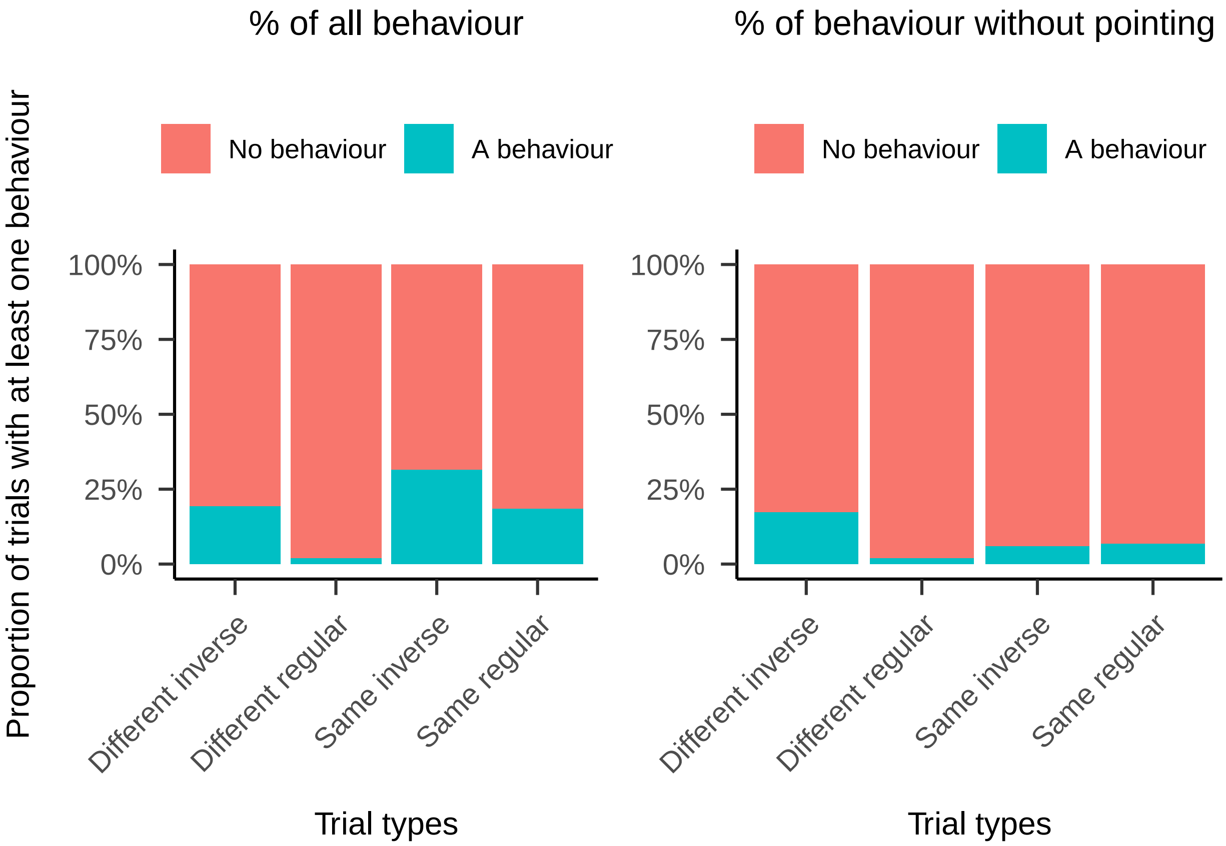

periment 2.

Testing sheets

Below, I include the code used to generate testing sheets for Experiment 1 as well as an example of a sheet (Figure B3.3).

```
order = data.frame(c('SP2', 'CO1', 'CO2'),
  sample(1:3), sample(1:3), sample(1:3), sample(1:3),
  sample(1:3), sample(1:3))
order = data.frame(b1=sample(1:3), b2=sample(1:3),
                   b3=sample(1:3), b4=sample(1:3),
                   b5=sample(1:3), b6=sample(1:3))
order[order==1] = 'SP2'
order[order==2] = 'CO1'
order[order==3] = 'CO2'
```

```
gaps = data.frame(g1 = sample(3:5, 6, replace = 1),
                  g2 = sample(2:5, 6, replace = 1),
                  g3 = sample(2:5, 6, replace = 1))


block1= c(rep.int(0, gaps[1,1]), order$b1[1],
          rep.int(0, gaps[1,2]), order$b1[2],
          rep.int(0, gaps[1,3]), order$b1[3])

block2 = c(rep.int(0, gaps[2,1]), order$b2[1],
           rep.int(0, gaps[2,2]), order$b2[2],
           rep.int(0, gaps[2,3]), order$b2[3])

block3 = c(rep.int(0, gaps[3,1]), order$b3[1],
           rep.int(0, gaps[3,2]), order$b3[2],
           rep.int(0, gaps[3,3]), order$b3[3])

block4 = c(rep.int(0, gaps[4,1]), order$b4[1],
           rep.int(0, gaps[4,2]), order$b4[2],
           rep.int(0, gaps[4,3]), order$b4[3])

block5 = c(rep.int(0, gaps[5,1]), order$b5[1],
           rep.int(0, gaps[5,2]), order$b5[2],
           rep.int(0, gaps[5,3]), order$b5[3])

block6 = c(rep.int(0, gaps[6,1]), order$b6[1],
           rep.int(0, gaps[6,2]), order$b6[2],
           rep.int(0, gaps[6,3]), order$b6[3])

#sheet = cbind(block1, block2, block3, block4, block5, block6)
```

Figure B3.3 Example of a testing sheet

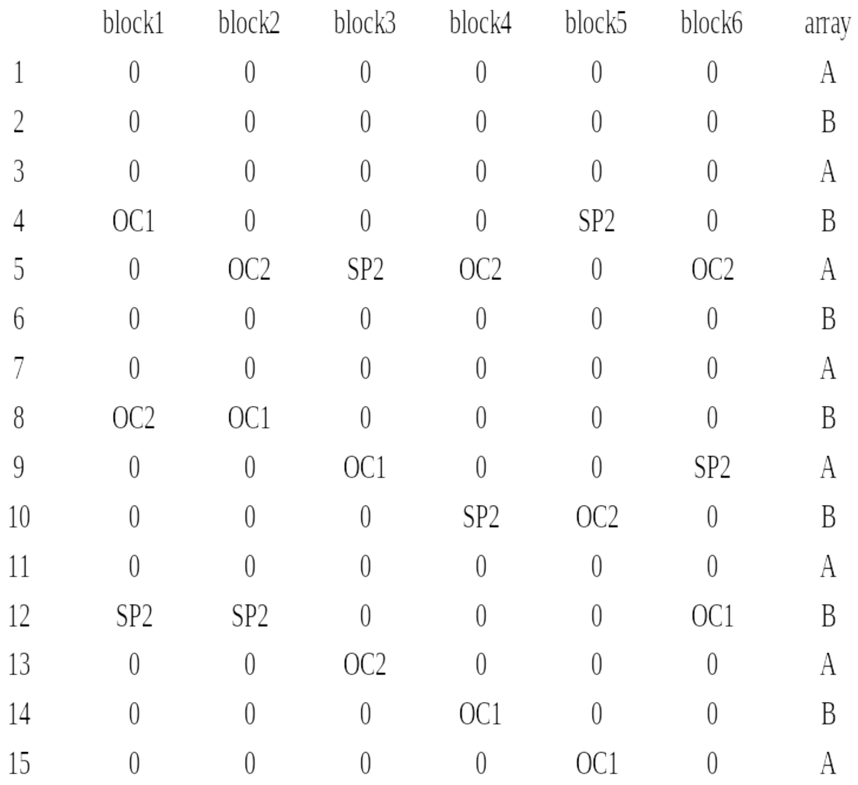

	block1	block2	block3	block4	block5	block6	array
1	0	0	0	0	0	0	A
2	0	0	0	0	0	0	B
3	0	0	0	0	0	0	A
4	OC1	0	0	0	SP2	0	B
5	0	OC2	SP2	OC2	0	OC2	A
6	0	0	0	0	0	0	B
7	0	0	0	0	0	0	A
8	OC2	OC1	0	0	0	0	B
9	0	0	OC1	0	0	SP2	A
10	0	0	0	SP2	OC2	0	B
11	0	0	0	0	0	0	A
12	SP2	SP2	0	0	0	OC1	B
13	0	0	OC2	0	0	0	A
14	0	0	0	OC1	0	0	B
15	0	0	0	0	OC1	0	A

Likewise, I include the code for generating and an example of testing sheets for Experiment 2.

```
#first define the type and number of trials in any given sessions; whatever
##you want in separate columns goes in different vectors

trials = c(rep('SAME_REG', 4), rep('DIF_REG', 4), rep('SAME_INV', 2), rep('DIF_INV',2))
leftright = c(rep('L', 6), rep('R', 6))
```

```r
whole = data.frame(c(1:12))

#right now this is set up to 4 sessions (length[whole]=[n_of_sessions*2]-1)
##this big loop combines the columns into a dataframe for any given participant
while (length(whole)<9){

overA= TRUE #this next loop generates the dif/same reg/inv trials
while (overA == 1) {
  block = sample(trials)
  a = rle(block)
  over2A = 3 %in% a$lengths
  over3A = 4 %in% a$lengths
  over4A = 5 %in% a$lengths
  over5A = 6 %in% a$lengths
  overA = sum(over2A, over3A, over4A, over5A) #no more than 2 subsequent sames
}

overB = TRUE #this next loop generates the left/right trials
while (overB == 1) {
  sides = sample(leftright)
  b = rle(sides)
  over3B = 4 %in% b$lengths
  over4B = 5 %in% b$lengths
  over5B = 6 %in% b$lengths
  overB = sum(over3B, over4B, over5B) #no more than 3 subsequent sames

}

whole = cbind(whole, sides, block)
}

whole = whole[,-1] # just removes the redundant order column (1:12)
```

Figure B3.4 Example of a testing sheet for Experiment 2.

sheet1

	s1	s2	s3	s4	
1	**L** SAME_REG	**R** DIF_REG	**R** SAME_REG	**L** DIF_INV	1
2	**L** SAME_REG	**L** DIF_REG	**L** DIF_REG	**L** DIF_REG	2
3	**R** DIF_INV	**R** SAME_REG	**R** SAME_INV	**R** SAME_INV	3
4	**L** DIF_REG	**L** SAME_INV	**L** DIF_REG	**L** DIF_REG	4
5	**L** DIF_REG	**L** DIF_INV	**R** DIF_INV	**R** DIF_INV	5
6	**L** SAME_REG	**R** SAME_REG	**L** DIF_REG	**L** SAME_REG	6
7	**R** DIF_REG	**R** SAME_REG	**L** SAME_REG	**R** DIF_REG	7
8	**R** SAME_INV	**L** DIF_REG	**R** DIF_REG	**L** SAME_REG	8
9	**L** DIF_INV	**R** SAME_REG	**R** DIF_INV	**R** SAME_INV	9
10	**R** DIF_REG	**L** SAME_INV	**L** SAME_REG	**R** DIF_REG	10
11	**R** SAME_INV	**L** DIF_REG	**L** SAME_INV	**L** SAME_REG	11
12	**R** SAME_REG	**R** DIF_INV	**R** SAME_REG	**R** SAME_REG	12

	s1	s2	s3	s4	
1	**L** DIF_REG	**R** DIF_INV	**L** DIF_REG	**L** DIF_REG	1
2	**R** DIF_INV	**R** DIF_REG	**L** SAME_INV	**L** SAME_REG	2
3	**R** DIF_REG	**L** SAME_REG	**R** SAME_REG	**R** DIF_REG	3
4	**L** SAME_INV	**R** SAME_REG	**L** SAME_REG	**R** DIF_REG	4
5	**R** SAME_REG	**R** DIF_REG	**R** DIF_REG	**L** SAME_INV	5
6	**L** SAME_INV	**L** SAME_INV	**L** SAME_REG	**R** SAME_INV	6
7	**R** DIF_REG	**L** DIF_REG	**L** DIF_REG	**L** SAME_REG	7
8	**R** SAME_REG	**R** DIF_INV	**R** SAME_INV	**L** SAME_REG	8
9	**L** SAME_REG	**L** SAME_REG	**L** SAME_REG	**L** DIF_INV	9
10	**L** DIF_INV	**L** SAME_INV	**R** DIF_REG	**R** SAME_REG	10
11	**R** SAME_REG	**L** SAME_REG	**R** DIF_INV	**R** DIF_REG	11
12	**L** DIF_REG	**R** DIF_REG	**R** DIF_INV	**R** DIF_INV	12

	s1	s2	s3	s4	
1	**L** DIF_REG	**R** SAME_INV	**L** DIF_REG	**L** DIF_INV	1
2	**R** DIF_INV	**L** SAME_REG	**L** SAME_REG	**L** DIF_INV	2
3	**R** SAME_REG	**R** DIF_REG	**R** SAME_REG	**R** DIF_REG	3
4	**R** SAME_INV	**R** DIF_INV	**R** DIF_INV	**R** DIF_REG	4
5	**L** SAME_INV	**L** SAME_REG	**L** DIF_REG	**R** SAME_REG	5
6	**L** SAME_REG	**R** SAME_REG	**L** DIF_REG	**L** SAME_REG	6
7	**L** DIF_REG	**R** DIF_INV	**R** SAME_INV	**R** SAME_INV	7
8	**R** DIF_INV	**L** SAME_REG	**R** SAME_REG	**R** SAME_INV	8
9	**R** DIF_REG	**R** DIF_REG	**L** DIF_INV	**R** DIF_REG	9
10	**R** SAME_REG	**L** DIF_REG	**L** DIF_REG	**L** SAME_REG	10
11	**L** DIF_REG	**L** SAME_INV	**R** SAME_INV	**L** DIF_REG	11
12	**L** SAME_REG	**L** DIF_REG	**R** SAME_REG	**L** SAME_REG	12

Page 1

Appendix C: Chapter 4

Code for generating testing sheets (Experiment 1)

```
num = c(1:12)
df = data.frame(num)

while (length(df)<5){ # tops it at 4 sessions

over = TRUE
while (over == TRUE){
  list = c(1:12)
  trials = c(rep('FAR_HQ', 3), rep('FAR_LQ', 3), rep('CLOSE_HQ', 3), rep('CLOSE_LQ', 3))
  block = sample(trials)
  subs = rle(block)
  over = 3 %in% subs$length  #this means that if there are 3 subsequents it'll run again, i.e., no more
  block = as.factor(block)
}

df = cbind(df, block)

}
```

Figure C4.1. Example of a testing sheet for experiment 1

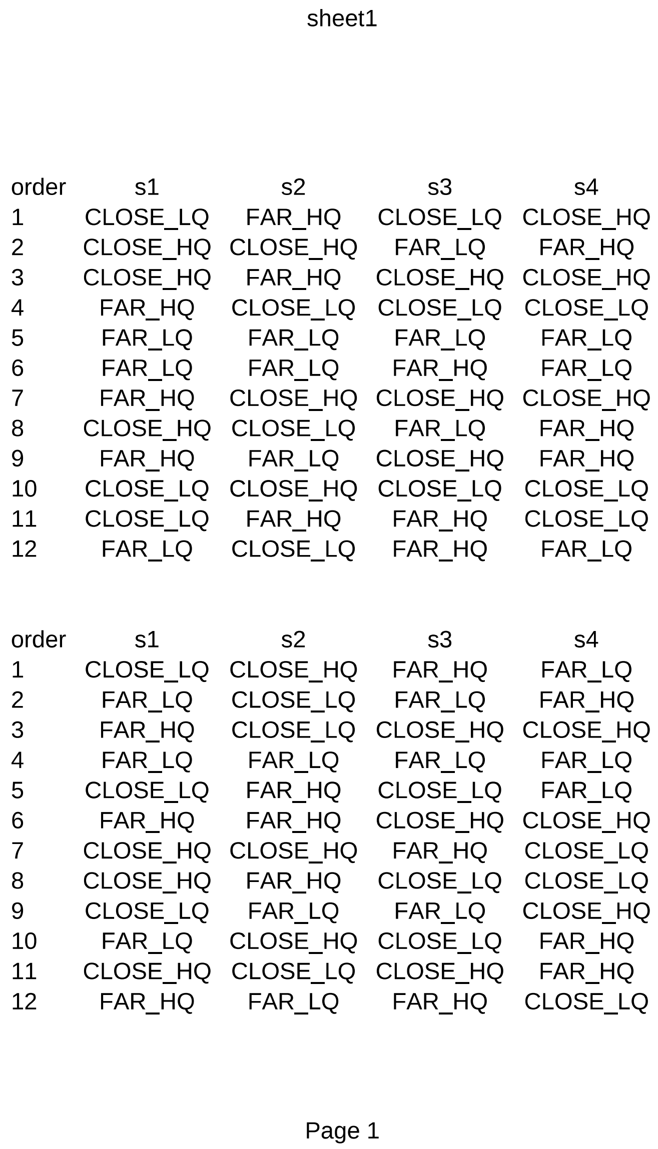

sheet1

order	s1	s2	s3	s4
1	CLOSE_LQ	FAR_HQ	CLOSE_LQ	CLOSE_HQ
2	CLOSE_HQ	CLOSE_HQ	FAR_LQ	FAR_HQ
3	CLOSE_HQ	FAR_HQ	CLOSE_HQ	CLOSE_HQ
4	FAR_HQ	CLOSE_LQ	CLOSE_LQ	CLOSE_LQ
5	FAR_LQ	FAR_LQ	FAR_LQ	FAR_LQ
6	FAR_LQ	FAR_LQ	FAR_HQ	FAR_LQ
7	FAR_HQ	CLOSE_HQ	CLOSE_HQ	CLOSE_HQ
8	CLOSE_HQ	CLOSE_LQ	FAR_LQ	FAR_HQ
9	FAR_HQ	FAR_LQ	CLOSE_HQ	FAR_HQ
10	CLOSE_LQ	CLOSE_HQ	CLOSE_LQ	CLOSE_LQ
11	CLOSE_LQ	FAR_HQ	FAR_HQ	CLOSE_LQ
12	FAR_LQ	CLOSE_LQ	FAR_HQ	FAR_LQ

order	s1	s2	s3	s4
1	CLOSE_LQ	CLOSE_HQ	FAR_HQ	FAR_LQ
2	FAR_LQ	CLOSE_LQ	FAR_LQ	FAR_HQ
3	FAR_HQ	CLOSE_LQ	CLOSE_HQ	CLOSE_HQ
4	FAR_LQ	FAR_LQ	FAR_LQ	FAR_LQ
5	CLOSE_LQ	FAR_HQ	CLOSE_LQ	FAR_LQ
6	FAR_HQ	FAR_HQ	CLOSE_HQ	CLOSE_HQ
7	CLOSE_HQ	CLOSE_HQ	FAR_HQ	CLOSE_LQ
8	CLOSE_HQ	FAR_HQ	CLOSE_LQ	CLOSE_LQ
9	CLOSE_LQ	FAR_LQ	FAR_LQ	CLOSE_HQ
10	FAR_LQ	CLOSE_HQ	CLOSE_LQ	FAR_HQ
11	CLOSE_HQ	CLOSE_LQ	CLOSE_HQ	FAR_HQ
12	FAR_HQ	FAR_LQ	FAR_HQ	CLOSE_LQ

Page 1

Code for generating testing sheets (Experiment 2):

```
trials = c(rep('HQF',6), rep('NULF', 6))

x = data.frame(c(rep('', 12)))



dfno2 = data.frame(c(1:12))
```

```r
while(length(dfno2)<25){
  over = TRUE
  while(over >= 1){
  block = sample(trials)
  a = rle(block)
  over2 = 3 %in% a$lengths
  over3 = 4 %in% a$lengths
  over4 = 5 %in% a$lengths
  over5 = 6 %in% a$lengths
  over = sum(over2, over3, over4, over5)
  }


  dfno2 = cbind(dfno2, block, x)


}
dfno2 = dfno2[,-1]
```

Figure C4.2. Example of a testing sheet for Experiment 2.

12 sheets tauzin 2 copy

	Josie		Tuli		
	s1	s2	s1	s2	
1	HQF	HQF	HQF	HQF	1
2	NO	HQF	NO	NO	2
3	NO	NO	HQF	HQF	3
4	HQF	HQF	HQF	HQF	4
5	NO	HQF	NO	NO	5
6	NO	NO	HQF	NO	6
7	HQF	HQF	NO	HQF	7
8	NO	NO	NO	NO	8
9	HQF	NO	HQF	HQF	9
10	HQF	HQF	HQF	HQF	10
11	NO	NO	NO	NO	11
12	NO	NO	NO	NO	12

	Flynn		Sam		
	s1	s2	s1	s2	
1	NO	NO	HQF	HQF	1
2	NO	HQF	HQF	NO	2
3	HQF	HQF	NO	HQF	3
4	HQF	NO	HQF	HQF	4
5	NO	HQF	NO	NO	5
6	HQF	NO	NO	NO	6
7	NO	HQF	HQF	HQF	7
8	HQF	NO	NO	NO	8
9	HQF	HQF	NO	HQF	9
10	NO	NO	HQF	HQF	10
11	NO	NO	NO	NO	11
12	HQF	HQF	HQF	NO	12

	Coco		Vicky		
	s1	s2	s1	s2	
1	NO	HQF	HQF	HQF	1
2	HQF	NO	NO	NO	2
3	HQF	NO	NO	NO	3
4	NO	HQF	HQF	HQF	4
5	HQF	HQF	HQF	NO	5
6	NO	NO	NO	NO	6
7	HQF	HQF	HQF	HQF	7
8	NO	HQF	NO	NO	8
9	NO	NO	HQF	HQF	9
10	HQF	HQF	NO	NO	10
11	NO	NO	NO	HQF	11
12	HQF	NO	HQF	HQF	12

Appendix D: Ethics approvals

University of St Andrews

School of Psychology & Neuroscience Ethics Committee

22 April 2021

Dear Kresimir

Thank you for submitting your application which was considered at the Psychology & Neuroscience School Ethics Committee meeting on 15th April 2021. The following documents have been reviewed:

1. Animal Ethics Form
2. Copy of Travel Risk Assessment Form (Coronavirus Infection Risk)

Project Title:	Appearance-reality discrimination in chimpanzees		
Researcher:	Kresimir Durdevic		
Supervisor:	Professor Josep Call		
Approved on:	15/04/2021	**Approval Expiry:**	15/04/2026
SEC Approval Code:	206		

The School of Psychology & Neuroscience Ethics Committee approves this study from an ethical point of view.

Approval is given for five years. Projects, which have not commenced within two years of original approval, must be re-submitted to the School Ethics Committee.

You must inform the School Ethics Committee when the research has been completed. If you are unable to complete your research within the five-year validation period, you will be required to write to the School Ethics Committee to request an extension or you will need to re-apply.

Any serious adverse events or significant change which occurs in connection with this study and/or which may alter its ethical consideration, must be reported immediately to the School Ethics Committee, and an application for ethical amendment Form submitted where appropriate.

Approval is given on the understanding that the ***ASAB Guidelines for the treatment of animals in behavioural research and teaching (ANIMAL BEHAVIOUR, 2018, 135, I-X)*** are adhered to.

Yours sincerely

Convenor of the School Ethics Committee

Ccs Professor Josep Call
Dr Tamara Lawson (Home Office Liaison Officer)

University of St Andrews

School of Psychology & Neuroscience Ethics Committee

22 February 2022

Dear Kreso

Thank you for submitting your application which was considered at the Psychology & Neuroscience School Ethics Committee meeting on the 3rd February 2022. The following documents have been reviewed:

1. Animal Ethics Form
2. External Permissions
3. Other

Project Title:	Mental State		
Researchers:	Elizabeth Warren, Kresimir Durdevic		
Supervisor:	Professor Josep Call		
Approved on:	22/02/2022	**Approval Expiry:**	22/02/2027
SEC Approval Code:	PS16011		

The School of Psychology & Neuroscience Ethics Committee approves this study from an ethical point of view.

Approval is given for five years. Projects, which have not commenced within two years of original approval, must be re-submitted to the School Ethics Committee.

You must inform the School Ethics Committee when the research has been completed. If you are unable to complete your research within the five-year validation period, you will be required to write to the School Ethics Committee to request an extension or you will need to re-apply.

Any serious adverse events or significant change which occurs in connection with this study and/or which may alter its ethical consideration, must be reported immediately to the School Ethics Committee, and an application for ethical amendment Form submitted where appropriate.

Approval is given on the understanding that the ***ASAB Guidelines for the treatment of animals in behavioural research and teaching (ANIMAL BEHAVIOUR, 2018, 135, I-X)*** are adhered to.

Yours sincerely

Dr Helen Sunderland
Administrator to the School Ethics Committee

Ccs Professor Josep Call
Home Office Liaison Officer