ORIGINAL ARTICLE

WILEY

# Broad-spectrum conceptual engineering

## Manuel Gustavo Isaac[1,2] iD

[1]Swiss National Science Foundation, Wildhainweg 3, Bern, 3001, Switzerland

[2]Arché Philosophical Research Centre, University of St Andrews, 17-19 College Street, St Andrews, Fife KY16 9AL, Scotland

**Correspondence**
Manuel Gustavo Isaac, Arché Philosophical Research Centre, University of St Andrews, 17-19 College Street, St Andrews, Fife KY16 9AL, Scotland.
Email: isaac.manuelgustavo@gmail.com

## Abstract

Conceptual engineering is the method for assessing and improving our representational devices. On its 'broad-spectrum' version, it is expected to be appropriately applicable to any of our representation-involving cognitive activities, with major consequences for our whole cognitive life. This paper is about the theoretical foundations of conceptual engineering thus characterised. With a view to ensuring the actionability of conceptual engineering as a broad-spectrum method, it addresses the issue of how best to construe the subject matter of conceptual engineering and successively defends the theses that conceptual engineering should be: (i) About concepts, (ii) psychologically theorised, (iii) as multiply realised functional kinds. Thereby, I claim to theoretically secure and justify the maximum scope, flexibility, and impact for the method of conceptual engineering on our representational devices in our whole cognitive life—in other words, a broad-spectrum version of conceptual engineering.

**KEYWORDS**

conceptual engineering, metaphilosophy, philosophical methods, theories of concepts

## 1 | INTRODUCTION

Conceptual engineering is the method for assessing and improving our representational devices. On its 'broad-spectrum' version, it is expected to be appropriately applicable to any of our representation-involving cognitive activities, with major consequences for our whole cognitive life. This paper is about the theoretical

foundations of conceptual engineering thus characterised—that is, about the framework for theorising the foundations of conceptual engineering as a method with a wide scope, a high flexibility, and a strong impact on our representational devices in our cognitive life. With a view to ensuring the actionability of conceptual engineering as a such broad-spectrum method,[1] the main problem this paper addresses is that of the subject matter of conceptual engineering, namely: What is the proper object domain of conceptual engineering? This problem in turn is subdivided into two successive issues. Firstly, how should we construe the representational devices that conceptual engineering target? Or more explicitly, should conceptual engineering target concepts? Secondly, if conceptual engineering is to target concepts, then, how should we theorise concepts for the very purposes of conceptual engineering? The thesis that I will defend is that conceptual engineering *should* target concepts, psychologically theorised, as multiply realised functional kinds. As I will argue at length, this option best theoretically secures and justifies the 'broad-spectrum' version of conceptual engineering—that is, as a method with a wide scope, a high flexibility, and a strong impact on our representational devices in our cognitive life.

The paper goes as follows: I start by providing a regimented overview of conceptual engineering as a metaphilosophical movement fitted with broad-spectrum ambitions (Section 2). Next, I address what I take to be the main foundational issue in conceptual engineering (namely, that of conceptual engineering's target domain) and solve it by establishing that conceptual engineering should target concepts, psychologically theorised (Section 3). Then, I develop further my proposal of a psychological conception of concepts for the purposes of conceptual engineering, that is, as multiply realised functional kinds (Section 4). Finally, I conclude with its main upshot—namely, that a psychological take on concepts as multiply realised functional kinds ensures the maximum scope, flexibility, and impact for the method of conceptual engineering on our representational devices in our cognitive life (Section 5).

## 2 | GENERAL BACKGROUND

This introductory section provides a quick regimented overview of conceptual engineering and further substantiates its broad-spectrum version. The purpose is to set the scene for the formulation of the main foundational issue that currently undermines the actionability prospects for its method (namely, that of its target domain), with a view to next solving that issue (namely, in the terms of a psychological conception of concepts). At first glance, conceptual engineering is a movement in metaphilosophy that promotes a renewed take on philosophical methodology as the study of our representational devices. Although there is no common framework unifying this renewed approach in the form of a research program, current work in conceptual engineering shares several core commitments.

### 2.1 | The core commitments

First, as concerns their theoretical framework, conceptual engineers typically take our representational apparatus to be made of cognitive devices that serve cognitive functions in the execution of certain cognitive tasks. One may think, for instance, of the 'four food groups' used to roughly classify food on an oversimplified nutritional basis (Nado, 2021, p. 15). Yet, basically any representational device we use when we want to understand,

---

[1]In what follows, 'actionability' means capability of being put to work and concerns the whole engineering process and its successive stages—that is, description, evaluation, design, test, and implementation (cf. Andow, 2020; Burgess & Plunkett, 2020; Cappelen & Plunkett, 2020; Chalmers, 2020; Eklund, 2021). The challenge it poses is thus larger, and includes the so-called 'implementation challenge' (Cappelen & Plunkett, 2020; Deutsch, 2020a, 2021; Jorem, 2021; Koch, 2021), which only deals with the last stage of the engineering process and concerns the issue of how to implement on the ground the ameliorative strategies that conceptual engineers may have come to advocate for.

explain, communicate, etc. something about something to someone or for oneself qualifies alike. Second, the main assumption that is then made on this basis is that the quality of our representational devices crucially determines that of the cognitive activity we use them for, whether practical or theoretical (thinking, talking, reasoning, behaving, etc.). For instance, one may take *phoneme* to be better than *letter*, which was previously in use, in order to well describe the sound pattern of a natural language, or *racialization* to be better than *race* in order to prevent discriminatory practices against socially formed racialized groups of people, and so on. To summarise, in a slogan: 'The better our representations are, the better our cognition will be. Third and as a corollary, the quality of our representational devices is then typically measured in terms of cognitive efficacy—that is, with regard to whether and how well a given representational device fulfils its cognitive function. Call what these three theoretical commitments of conceptual engineering together form its 'cognitive functionalism.'

Next, as regards the practical application of conceptual engineering, it is typically expected among fellow engineers that our representational devices can be (re-)engineered and for the better—call this their 'ameliorative interventionism.'[2] One key feature of the practice of conceptual engineering is then its normativity: Conceptual engineers aim to prescribe what representational devices we should have and use, instead of merely describing those we do already have and use as a matter of fact. With this in mind, the application of conceptual engineering as a method for carrying out ameliorative projects on specific case studies commonly follows a two-step process. First is an assessment phase in which the representational device under consideration is functionally assessed for what it does, or what we do with it, and for how well it does it, with a view to identifying some improvable features in it (such as deficiencies to fix). Second is an improvement phase in which strategies for improving the identified improvable features of the representational device under consideration (such as for fixing its deficiencies) are functionally prescribed, with a view to determining whether and how we should use a given representational device. Importantly, these two successive phases both pertain to the normative dimension of conceptual engineering and must be supplemented by some preliminary descriptive work that provides them with the background data on which to operate.

Against the background of its theoretical and methodological commitments—namely, its 'cognitive functionalism' and 'ameliorative interventionism'—conceptual engineering can then be characterised in a nutshell as a metaphilosophical movement that aims to normatively ameliorate the cognitive quality of our representational devices in terms of functional efficacy. That is the characterisation we will use in the remainder of this paper, in particular, when adjudicating between the different possible target objects for conceptual engineering (Section 3).

## 2.2 | The main desiderata

With this characterisation in mind, we may then further substantiate the idea of a broad-spectrum version of conceptual engineering. To this end, we can introduce three parameters that help measuring the strength of the method of conceptual engineering as regards its application to our representational devices, namely: Impact, scope, and flexibility. Impact concerns the type of connection conceptual engineering has to people's mind (direct vs. indirect); scope, the range of its tractable representational devices (wide vs. narrow); and flexibility, the appropriateness of its calibration (fine- vs. coarse-grained). A broad-spectrum version of conceptual engineering is then one that aims for the application of its method to have a wide scope, a high flexibility, and a strong impact on our representational devices in our cognitive life.

Impact, scope, and flexibility together underlie the metaphilosophical movement of conceptual engineering as its common desiderata, whether overtly or not. As concerns scope and flexibility, their rationale lies in the 'Master

---

[2]Cappelen's (2018, pp. 72–74) 'lack of control' principle may be the exception here, as it somehow dooms the implementation of representational changes to remain infeasible.

Argument,' which implicitly underwrites the practical application of conceptual engineering.[3] The argument can be reconstructed and generalised as follows:

P. 1 Contingency: The present state of our representational devices is the product of contingencies—that is, they could have been otherwise;

P. 2 Suboptimality: In their contingent present state, our representational devices are unlikely to be optimal, with regard to the function they are meant to fulfil;

C. 1 Representational skepticism: We should critically assess the representational devices we use in our cognitive activities, with regard to the function they are meant to fulfil[4];

P. 3 Cognitive detriment: The suboptimal fitness of our representational devices to the function they are meant to fulfil is detrimental to our cognitive behaviours and abilities (Section 2.1);

C. 2 Ameliorative interventionism: We should aim to improve the representational devices we use in our cognitive activities, with regard to the function they are meant to fulfil.

Following the Master Argument thus reconstructed, conceptual engineering is de facto required to be appropriately applicable to any of our representation-involving cognitive activities—hence, the wide scope and high flexibility as its common desiderata.

As for impact, it can in turn be motivated by the following basic observation: Representational amelioration is never sought for its own sake by conceptual engineers. Rather, it is nearly always taken as a means to bring about broader changes in our cognitive life—for instance, to combat systemic social oppression, to dismantle misleading systems of beliefs, to formulate generalisations of explanatory value, or to promote hermeneutic justice among individuals. And given the critical role our representational devices play in our cognitive economy, the consequences of ameliorative changes to our representational schemes and repertoires for our whole cognitive life are expected to be major. As Burgess and Plunkett (2013, pp. 1091–1092) have observed:

> Our [representational] repertoire determines not only what we can think and say but also, as a result, what we can do and who we can be. In other words, which [representations] we use has an important impact on the space of possible actions and lives available to us.

Hence, strong impact as another of conceptual engineering's common desideratum.

All in all, a broad-spectrum version of conceptual engineering can then be specified as one that aims for its method to be appropriately applicable to any of our representation-involving cognitive activities (wide scope and high flexibility), with major consequences on our cognitive life (strong impact).[5] That is the other characterisation we will use in the remainder of this paper, both when adjudicating between the different possible target objects for conceptual engineering (Section 3) and when developing one of its specific variants (Section 4), in order to make it an actionable method.

---

[3]'The Master Argument' gives its title to Cappelen (2020), where it applies to the relationship between words and their meanings, with a specific focus on 'core philosophical terms' (Cappelen, 2020, p. 134) and no explicit connection to the so-called 'representational skepticism' (see below); the argument is also implicitly at work in genealogical/historical approaches to conceptual engineering (e.g., Dutilh Novaes, 2016, 2020; Plunkett, 2016).

[4]'Representational skepticism' is introduced by Cappelen (2018, pp. 5–7) as the 'intellectual attitude or cognitive disposition' that is distinctive of conceptual engineers, who never 'uncritically take over the representational devices handed to them' (Cappelen, 2018, p. 5).

[5]By contrast, a narrow-spectrum version of conceptual engineering is one whose consequences for our cognitive life are minor (little impact) (for instance, confined to the academic classroom), or whose application either is restricted to some of our representation-involving activities only (limited scope) (for instance, our theoretical activities), or obeys the same unique standards wholesale (low flexibility) (for instance, those of scientific rationality).

# 3 | THE MAIN FOUNDATIONAL ISSUE

Making conceptual engineering an actionable method requires us to theorise its target domain. This is because a method that lacks any grasp on what it is about is presumably bound to remain a useless piecemeal practice. As Cappelen (2018, p. 141) asserted: '[Y]ou don't really have an account of conceptual engineering unless you make an explicit choice here.' The purpose of this core section is thus to tackle the problem of conceptual engineering's target domain, which currently undermines the prospects for making it an actionable method—call it 'the main foundational issue.' With this in mind, we will proceed in two steps: First, we will reflect on how best to construe the representational devices that conceptual engineering target; second, how best to theorise concepts for conceptual engineering, if these prove to be what conceptual engineering should be targeting. The solution proposed to the main foundational issue is then expected to turn conceptual engineering into a well-founded method, improving thereby its actionability prospects, especially when taken on its broad-spectrum version.

## 3.1 | Which representational devices?

The strategy for tackling the issue of which representational devices conceptual engineering should work on is pretty simple: I start with a taxonomy of the different possible objects of conceptual engineering and then select the most conducive one with a view to making conceptual engineering an actionable method. This strategy is meant to help us progressively narrow down the target domain of conceptual engineering, so as to next identify the framework that is best suited for theorising the foundations of conceptual engineering.

### 3.1.1 | A basic taxonomy

Several different positions have been taken so far vis-à-vis the target domain of conceptual engineering. These can be taxonomised via a twofold alternative made of two successive dichotomies. The first alternative is a general one that consists of two options: On the one hand, the principled views on which conceptual engineering is taken to target some specified representational devices (viz., concepts, linguistic meanings, lexical items, subject matters, etc.); and on the other hand, the unprincipled views on which conceptual engineering may well target any kind of representational device more or less indistinctly, that is, with no real need of further specification (e.g., Burgess, 2020; Burgess & Plunkett, 2013, 2020; Cantalamessa, 2021; Nado, 2021; Prinzing, 2018; Sterken, 2020; Tanswell, 2018).

The second alternative is a specific one that distinguishes, within the principled views, between the two following options: A positive wing, on which the specified representational devices that conceptual engineering target are to be construed as concepts (e.g., Haslanger, 2020a, 2020b; Isaac, 2020; Koch, 2020; Machery, 2017; Richard, 2020; Scharp, 2013, 2020) and a negative wing, on which these representational devices are not to be construed as concepts, but instead, for instance, as conceptions, linguistic meanings, intensions, or more generally, as the ways in which we think and talk about topics (e.g., Cappelen, 2018; Sawyer, 2018, 2020a, 2020b). Now, I contend that, among these three options (that is, the unprincipled views along with the positive and the negative principled views), only one best suits the purpose of making conceptual engineering an actionable method—namely, the pro-concept one.

### 3.1.2 | The comparative selection

The argument that concepts are the proper representational devices for conceptual engineering to target proceeds in two steps. The first step concerns the general dichotomy between the principled and unprincipled views on the target domain of conceptual engineering. It consists in establishing that the representational devices which that

target domain is made up of should be specified, as opposed to being left unspecified. Drawing on common sense, the argument to this effect is the following: Conceptual engineering is first and foremost devised as a method, and a method is always better off grasping its object. More precisely, as a way of doing something to attain a specific object (*meta-odos*), a method consists of a collection of principles and/or procedures, which themselves serve a given field of study.[6] And since the field of study in question supplies these principles and/or procedures with the target objects they are operating on, a method that does build on a theory of its target domain is always more likely to be efficiently and consistently actionable. Therefore, in order to make conceptual engineering an actionable method, one should start by specifying what its target domain should be. Otherwise, if one were to disregard the need to specify the target domain of conceptual engineering, one would then most certainly condemn conceptual engineering to remain a useless, piecemeal collection of resembling practices, as it will then intrinsically lack any overall grip over its target domain. The underspecification of the target domain of conceptual engineering would thereby hinder the prospect of making conceptual engineering an actionable method, which, in light of its 'ameliorative interventionism' (Section 2.1), is among the central proposes of the whole enterprise. In other words, the un-principled views on the target domain of conceptual engineering deliver a pragmatically inconsistent option.

Next, once we acknowledge that the target domain of conceptual engineering should be specified in terms of representational devices, the second step of the argument to establish that these should be construed as concepts, instead of any other kind of representational devices, concerns the specific dichotomy that distinguishes, within the principled views on the target domain of conceptual engineering—the pro-concept and the no-concept options. The argument to this effect is pretty straightforward, as it proceeds from a *reductio* that goes as follows: If conceptual engineering is not (to be) about concepts, then 'conceptual engineering' is a misnomer (Cappelen, 2018, pp. 53, 104).[7] But if 'conceptual engineering' is a misnomer, then it is misleading. And if it is misleading, then it constitutes a cognitive impediment—typically, it will trigger wrong associations in people's mind. Therefore, the very label of conceptual engineering would contravene the core rationale of conceptual engineering, that is: Creating better representational devices to foster better cognitive activities (thinking, talking, reasoning, etc.). As Cappelen (2018, p. 53 <ms.>) concedes, on the no-concept wing, '"conceptual engineering" isn't a great label <(or even a very bad one [...])>.' And this would then confront conceptual engineers with a self-discrediting predicament. In light of conceptual engineering's 'cognitive functionalism' (Section 2.1), 'conceptual engineering without concepts [only]' (Greenough, 2020) thus appears to deliver yet another pragmatically inconsistent option, while the pro-concept wing prevails.[8]

A third option so far ignored within the principled views would be to explicitly advocate a form of ecumenism toward our representational devices by including them all as possible targets for conceptual engineering projects.[9] However, this option faces the following dilemma: If it does not specify further its target objects, conceptual engineering will then most certainly remain a piecemeal collection of merely resembling practices that lack any overall grip over its target domain (like the unprincipled views); on the other hand, if it specifies its target objects further in terms of, e.g., intensions, conceptions, speaker-meanings, lexical items, concepts, etc.,

---

[6]See 'method' in the *OED*: 'The principles or procedures of any mode or field of cognitive activity, themselves considered as an object or branch of study' (entry I.2d); 'A special form of procedure or characteristic set of procedures employed (more or less systematically) in an intellectual discipline or field of study as a mode of investigation and inquiry' (entry I.3a).

[7]Cf. Cappelen (2018, p. 104): 'Conceptual engineering and amelioration are mislabeled: it's not about improving concepts—in fact, it's not about concepts at all.'

[8]One might think that 'conceptual engineering' should be treated as a name, rather than a description, which simply refers to 'the activity that self-described [sic] conceptual engineers are engaged in' (Cappelen, 2018, pp. 3, 4), but the problem with this proposal is that it turns conceptual engineers into 'exploiters,' who exploit the (misleading) 'lexical effects' associated with their label to manipulate people's expectations (Cappelen, 2018, pp. 133–134).

[9]Target domain ecumenism in conceptual engineering is best illustrated here: 'We would like to cast as wide a net as possible. Eliminativists about concepts will hopefully be able to massage our discussion to fit the mold of their favorite metaphysics of mental representation. [...] Theorists with diverse views on the nature of content ought to be able to engage in conceptual [engineering] without talking past each other.' (Burgess & Plunkett, 2013, p. 1095).

conceptual engineering will then most certainly include contradictions among the principles guiding its practical application to specific case studies. In either case, an ecumenical approach to the target objects of conceptual engineering would thus hinder the prospect of making it an actionable method. Finally, such an ecumenical approach would also fall prey to the self-discrediting predicament for conceptual engineering (see above). Therefore, conceptual engineering should target concepts and concepts only.[10]

## 3.2 | Which theory of concepts

Once we have established that the representational devices that conceptual engineering target should be construed as concepts, on pain of pragmatic inconsistencies otherwise, the next step to address conceptual engineering's target domain problem is to answer the question of how best to theorise concepts for the purposes of conceptual engineering—against the laid-back attitude that is widespread among pro-concept engineers.[11] Here as above, my strategy is pretty simple: I start again with a taxonomy of the most important types of theories of concepts that are available in the literature, and then I select the most conducive one with a view to securing a broad-spectrum version of conceptual engineering. As already stated, this will eventually provide us with a baseline from which to develop further the concept of concept at work in the chosen theoretical framework for the purposes of conceptual engineering.

### 3.2.1 | A basic taxonomy

The taxonomy for distinguishing the different theories concepts that are available to conceptual engineering is pretty basic, again, in that it builds on an alternative between two main frameworks, which further correlate with different models of and approaches to concepts (Johnston & Leslie, 2012; Löhr, 2020; Machery, 2009; Peacocke, 1992; Rey, 1983, 1985). The first framework available for theorising concepts is that of mainstream analytic philosophy of mind and language (e.g., Burge, 1993; Fodor, 1975; Peacocke, 1992). It is thus the one that is commonly adopted, explicitly or not, by the 'positive wing' in conceptual engineering, on which conceptual engineering is taken to target concepts (e.g., Brun, 2016; Burgess & Plunkett, 2013; Chalmers, 2011; Plunkett, 2015, 2016; Prinzing, 2018; Simion, 2018; Thomasson, 2020). In this framework, concepts are construed as semantic entities, further endowed with a 'semantic structure' (Margolis & Laurence, 2010, p. 219). As such, they are typically modelled as the sub-components of our thought contents, which in particular serve a referential function (e.g., Burge, 1993). Concepts are thus approached, in this framework, with a focus on their semantic constituency. And the goal here has classically been to deliver 'a priori, analytic truths about the world' (Machery, 2017, p. 209).[12]

The alternative framework for theorising concepts that is available to conceptual engineering is that of the psychology of concepts. By contrast to the philosophical framework, only a few attempts have been made so

---

[10]Importantly, however, the pragmatic inconsistency of anything but the pro-concept option for conceptual engineering does not preclude the possible pluralisation of engineering projects, which should then be labeled 'semantic', 'lexical,' 'terminological engineering,' and so on, depending on what representational devices their target objects are taken to be (linguistic meanings, lexical items or terms, etc.).

[11]As observed by Cappelen (2018, p. 141): 'There's of course already a smorgasbord of options for how to think about concepts. [...] However, and this is the strange part, those who talk of conceptual engineering as operating on concepts don't start by making choices on this smorgasbord. They often just talk about "concepts", their engineering, and then leave it at that.' And tentatively explained by Deutsch (2020a, p. 4): 'Proponents of th[e] standard account of conceptual engineering tend to say very little about what concepts are. I suspect this is because they don't think much needs to be said about this. Philosophers speak of concepts all the time [...]. Proponents of the standard account of conceptual engineering appear to think that there is no problem simply co-opting this "concept talk", trusting that philosophers, at least, will know how to interpret it.'

[12]This generic characterisation of philosophical theories of concepts is neutral with regard to their further specification, first, as dealing with concepts as abstract objects or mental representations (or as a mixture of both) (Margolis & Laurence, 2007), second, as dealing with conceptual structures based on either a 'Containment Model,' as per classical definitionism (among others), or on an 'Inferential Model' (Margolis & Laurence, 1999), and third, as building on an internalist or instead on an externalist theory of mind and language.

far to capitalise on this research field for the purposes of conceptual engineering (Isaac, 2020; Koch, 2020; Machery, 2017). In this framework, concepts are construed as cognitive entities, further endowed with a 'processing structure '(Margolis & Laurence, 2010, p. 219). As such, they are typically modelled as the *explanans* of our higher cognition, which in particular serve categorisation functions. Concepts are thus approached, in this framework, with a focus on their cognitive efficacy. And the goal is here to deliver 'empirical propositions about the mind' of descriptive and/or explanatory value (Machery, 2017, p. 209). Now that we have distinguished between these two different types of framework for theorising concepts, let us comparatively assess them with regard to their conduciveness to securing a broad-spectrum version of conceptual engineering.

### 3.2.2 | The comparative selection

In order to assess whether the philosophical or the psychological framework for theorising concepts is the most conducive to conceptual engineering, I will use as criteria the three parameters that make it a broad-spectrum method—namely, impact, scope, and flexibility. Against this backdrop, I will contend that, in all of these respects, the psychological framework for theorising concepts outclasses the philosophical one for the purpose of conceptual engineering, on its broad-spectrum version.

First, as concerns impact, the issue at stake is whether the framework under consideration allows for a direct connection or not with people's minds and cognitive life, abilities, and behaviours. The reason why the psychological framework fares better than the philosophical one in terms of impact thus characterised is simply that, while being taken as cognitive entities, concepts operate as the constitutive principles for the cognitive processes and competences of real cognitive agents: '[They] are the tracks our minds prefer to travel on' (Machery, 2017, p. 222). Consequently, the psychological concept of concept turns conceptual engineering into an operation of changing or fixing people's minds and ways of thinking when these appear to be improvable. Thereby, it ensures a strong impact for the method of conceptual engineering on our cognitive life. By contrast, on the philosophical framework, while concepts are taken as semantic entities, they act as mere regulatory principles for the epistemic (viz., knowledge-related) processes and competences of some ideal epistemic agent. This becomes clear in light of what arguably are the two most important desiderata that bear on any philosophical theory of concepts, along with their corresponding explanatory purposes. These are, first, the stability of conceptual content, which serves to explain two phenomena: The agent's rational/logical inference-making, on the one hand, which is both an idealisation on the agent's actual cognitive competences and performance and pertains in itself to the epistemic realm; and on the other hand, the 'publicity' of conceptual contents, which proceeds from an idealisation on the agent's actual conceptual content and is typically endowed with some epistemic normativity as regards how agents actually have and use them (in other words, they can be wrong). The second desideratum is the compositionality of conceptual content, which serves to explain the systematicity and the productivity of the agent's thought, both of which properties are nowadays quite commonly taken to be sheer idealisations. In this light, philosophical theories of concepts lead us to formulate possession and application conditions of concepts that are primarily met, if not only, by agents when these are both idealised and epistemic.[13] Consequently, the philosophical concept of concept threatens to turn conceptual engineering into the rather frivolous activity of fiddling with representational contents whose bearing on people's cognitive abilities and behaviours remains rather dubious, to say the least. Thereby, it allows the method of conceptual engineering to have only a limited impact on our actual cognitive life.

---

[13]Note that psychological theories of concepts also need conceptual contents to be stable and to compose, but in a 'much more relaxed way' (Löhr, 2020, p. 12)—that is, in a way that fits the adaptation of our behaviours and abilities to the complexities of life in the real world.

Next, as concerns both scope and flexibility, the issue at stake is whether the framework under consideration allows for an exhaustive or a restrictive application of conceptual engineering to the diversity of our conceptual devices (cf. Section 2). The reason why the philosophical framework fares worse than the psychological one in terms of scope and flexibility thus characterised is that, while being taken as semantic entities, concepts are primarily driven by truth requirements and further constrained by epistemic standards. In essence, they mean to enter, as unsaturated components, the formation of full thought contents, which in turn ultimately aim at being true of their referents and thereby provide the agent with knowledge about these.[14] Yet as a consequence, in order to fulfil their alethic-epistemic agenda, engineering projects of the semantic type face the following dilemma: Either to focus on theoretical concepts, that is, on concepts of theoretical entities (the mathematical concept of probability, the semantic concept of truth, the astronomical concept of planet, etc.), and thereby restrict the scope of conceptual engineering on our conceptual devices; or to focus on theoretical purposes, and thereby reduce the flexibility of conceptual engineering to the same, unique standard (viz., that of scientific rationality). By contrast, on the psychological framework, while they are taken as cognitive entities, concepts are first and foremost constrained by pragmatic standards and ultimately driven by efficacy requirements. In essence, they serve to enable the execution of the cognitive tasks we use them for. Here, in Nado's (2021, pp. 17, 4) words: 'We could say, then, that the goal of conceptual engineering is to design effective concepts,' while 'truth and knowledge, by contrast, are secondary goals at best.' Consequently, in order to fulfil their pragmatic-functionalist agenda, engineering projects of the cognitive type need to maximally enlarge the scope of their tractable concepts (from the more mundane to the most sophisticated ones), for the sake of exhaustiveness, while respecting their specific identities and purposes as different concepts via the recourse to multiple, different standards, in a highly flexible way, for assessing and improving their cognitive efficacy.

Bottom-line: The psychological framework for theorising concepts outclasses the philosophical one not only in terms of impact, but also in terms of scope and flexibility for the method of conceptual engineering on our conceptual devices in our cognitive life. Therefore, it is the one that best secures a broad-spectrum version of conceptual engineering.

## 3.3 | Upshot

This section has tackled the target domain problem for conceptual engineering in two successive steps. First, I argued that concepts are the proper target objects for conceptual engineering, as opposed to any other kind of representational devices, on pain of otherwise being pragmatically inconsistent. Second, I argued that concepts for conceptual engineering should be psychologically theorised as cognitive entities, as opposed to their philosophical theorisation as semantic ones, on pain of otherwise limiting the spectrum of its method. The upshot is to turn conceptual engineering into a well-founded method, thus improving the prospect of making it actionable, especially when taken on its broad-spectrum version. In this respect, the framework of psychological theories of concepts provides us with the baseline from which to develop further the concept of concept so that it best suits the proposes of conceptual engineering, on its broad-spectrum version.

---

[14]In the literature on conceptual engineering, the primacy of epistemic responsiveness for concepts is clearly captured by Simion (2018, p. 923): 'Concepts, just like beliefs, are representational devices, their function is an epistemic one: to represent the world. In virtue of this function, concepts will be properly functioning when responsive to epistemic reasons, and malfunctional when responsive to practical reasons. Concepts will be good concepts qua concepts when they are epistemically good.'

# 4 | A POSITIVE PROPOSAL

The purpose of this final section is to develop further the concept of concept at work in the psychological framework for theorising concepts in order to improve its conduciveness to the purposes of conceptual engineering. This improvement will be conducted in line with two main desiderata: First, to fit well with the cognitive functionalism that characterises the theoretical framework of conceptual engineering, that is, with its take on our representational apparatuses as made up of cognitive devices; second, to support the broad-spectrum ambitions for the practical implementation of the method of conceptual engineering, that is, for its impactful and calibrated application to a diversity of tractable items that ranges from the more mundane to the most sophisticated concepts. These two desiderata will be satisfied by proposing a psychological characterisation of concepts as functional kinds, and then by further specifying them as multiply realised. Thereby, we will theoretically secure and justify a wide scope, a high flexibility, and a strong impact for conceptual engineering as a method on our conceptual devices in our cognitive life—in other words, a broad-spectrum version of conceptual engineering.

## 4.1 | Concepts as functional kinds

The way in which the notion of 'concept' is used in cognitive science allows for a consensual regimented characterisation (e.g., Machery, 2005, 2009, 2017). Concepts, as regards their content, are typically taken in cognitive science as bodies of information about some categories of referents (individuals, substances, classes, event-types), which form a stable core within people's belief-like states. The states which concepts are made up of are belief-like states, and not full-blown beliefs, since they are typically opaque (that is, not easily accessible to consciousness) and, most importantly, since they do not necessarily obey rationality constraints (Machery, 2017, p. 210, fn. 4). As a corollary, in contrast to the semantic approach, conceptual contents on their cognitive take are neither primarily, nor necessarily, driven by some truth criterion and/or responsive to epistemic standards—in the sense that they do not essentially have to provide us with true beliefs about the world. The alethic-epistemic framework is, for the most part, irrelevant, or secondary at best, for making sense of how they actually work in people's minds. Furthermore, concepts on the psychological picture are located within the individual agent only, which contrasts with the 'content publicity' (or 'sharability') requirement that bears upon the semantic take on concepts. Thus characterised, conceptual contents can then be specified in two steps that concern their status and function, respectively.

### 4.1.1 | By-default categorizing devices

First, with regard to their status, the belief-like states that constitute our conceptual contents are retrieved by default (and thus defeasibly) from our long-term memory, which is the memory that is accessed by the cognitive processes that underlie the exercise of our higher cognitive competences in the execution a higher cognitive task (more on this below). The default retrieval of concepts *qua* subsets of belief-like states can be delineated by the following homeostatic cluster of properties: Speed, they are rapidly retrieved; automaticity, their retrieval is uncontrollable; context-independence, they are retrieved in every context (Machery, 2015). The invariantism that supports this by-default characterisation of information retrieval for conceptual content explains why a substantial amount of retrieved information overlaps across different contexts in cases of concept-involving cognition. It thus accounts for the privileged status of concepts in the economy of our higher cognition—backing up thereby one of the most important, but by and large unsubstantiated, assumptions of conceptual engineering. Consequently, owing to its defaultness condition, which makes changes to our conceptual contents hard but not impossible to implement (more on this below), invariantism avoids the risk of turning conceptual engineering

into a trivial and innocuous belief revision project (cf. Deutsch, 2020a, 2020b; Pinder, 2019): On psychological invariantism, concepts are not whatever informational content you may fancy to coin on the fly, but those which are endowed with a central role in our higher cognitive economy in virtue of their default retrieval.[15]

Next, with regard to their function, conceptual contents cognitively characterised can be further specified as playing a causal and explanatory role in the cognitive processes that underlie the higher cognitive competences that are exercised in the execution of a higher cognitive task by a cognitive agent. The cognitive processes here in play consist of a series of operations that access the information stored in the agent's long-term memory in order to bring about the functions to be fulfilled by the exercise of a given higher cognitive competence. These operations in turn can be reduced to category-abstraction, on the one hand, which is 'the "bottom-up" process of extracting information from a single encounter with an object or property-instance, and generalising such information to all encounters with that object or property,' and category-induction, on the other hand, which is 'the complementary "top-down" process of projecting such knowledge to new encounters' (Lalumera, 2010, p. 218). As for the functions fulfilled by the higher cognitive competences, they are typically those of categorisation, deduction, induction, action-planning, analogy-making, and linguistic understanding. Finally, as I understand them, the higher cognitive tasks whose execution is achieved through the exercise of these higher cognitive competences simply correspond to the overall cognitive activity they contribute to performing in a given context (reading, listening, talking, reasoning, planning, etc.). In light of their function thus explicated, conceptual contents cognitively characterised can be seen as by-default 'categorizing devices' (Löhr, 2020).

## 4.1.2 | Refurbishing the mind

Concepts *qua* default categorizing devices most notably work by underwriting the mind's default inferences—that is, '[the] inferences the mind is disposed to draw, that, so to speak, spring to mind, that it only resists when attention is drawn to particular facts that defeat this disposition' (Machery, 2017, p. 222)[16]—yet in a way that is flexible enough to fit with the complexities of the world in which we live (cf. note 13). In effect, on the invariantism that underpins the by-default characterisation of concepts, the variation of the overall information retrieval across contexts in cases of concept-involving cognition is presumed to undergo a two-step process that ensures the context-dependence of its activation along with the task-appropriateness of its use (Machery, 2009, 2010, 2015). First, the default body of information that corresponds to a concept is invariantly retrieved from long-term memory; second, its deployment is then adapted to the specifics of a particular situation via its combination with some further context-dependent background information on the category of referents that the default body of information is about—for instance, contextual information may cancel the automatic deployment of stereotypical inferences (more on this below). This two-step process of informational retrieval for concepts is exactly how, why, and where conceptual engineering matters.

On their cognitive characterisation, concepts are functional kinds that determine the ways in which we think about things via the class of inferences that their default body of information underwrite. The overall process of conceptual engineering then amounts to, first, making explicit the typically opaque role 'concepts play in people's cognitive life' (Machery, 2017, p. 224), and second, either re-delineating our conceptual contents or mitigating their deployment when the inferences they underwrite are found improvable. The re-delineation of our conceptual contents would consist either in removing some of their composing elements or in including new elements in them, for instance, those which previously belonged to the background information—both of which, in turn, can be achieved by modifying the relevant parameters in the agent's environment with which

---

[15]See Machery (2015) for empirical evidences supporting invariantism over its contextualist alternative.

[16]Note that the class of concept-underwritten inferences is here not semantically or epistemically characterised, for '[it] neither articulates the semantic content of the concept (its meaning), nor does it have any distinctive justificatory status' (Machery, 2017, p. 222).

their default retrieval co-occurs. On the other hand, the mitigation of the inferential pattern of our conceptual contents would consist in blocking the automatic deployment of the stereotypical inferences they underwrite and replacing these by non-stereotypical ones—which, in turn, can be achieved by linguistic means, via the explicit marking of the deviation from stereotypes or from salient uses with contextual cues (Fischer, 2020). In short, on the psychological approach, conceptual engineering would thus amount to a mind-refurbishing process that aims to ameliorate the efficacy of concepts as cognitive devices and thereby essentially change the way in which people think about things so that it is better suited to executing the cognitive task it underpins. Besides fitting very well with the cognitive functionalism of conceptual engineering, this characterisation of concepts as functional kinds thereby theoretically secures and justifies the broad-spectrum ambitions of its research program when these are cashed out in terms of their impact on our conceptual devices for our whole cognitive life.

## 4.2 | Concepts as multiply realised

The next step is to theoretically secure and justify the broad-spectrum ambitions of conceptual engineering as a research program in terms of scope and flexibility on our conceptual devices in our cognitive life. To this end, I will further specify my proposal of a psychological characterisation of concepts as functional kinds in terms of multiple realisation.[17] The idea goes as follows: As is well known, three main theoretical paradigms have emerged in reaction to the failure and rejection of the classical theory of concepts—namely, the exemplar-, the prototype-, and the theory-theories of concepts.[18] All three main paradigms posit different theoretical entities that they respectively take concepts to be, that is: Exemplars, which are bodies of information about the properties of some single exemplar(s) of the concept's category of referents (individual members of a class, particular samples of a substance, or instances of an event-type) and which the other members of what the concept is about resemble (e.g., Medin & Schaffer, 1978; Smith & Medin, 1981); prototypes, which are bodies of statistical information about the properties that are typical and diagnostic of the members of the concept's category of referents (e.g., Rosch, 1973, 1975; Rosch & Mervis, 1975); and theories, which are structured bodies of explanatory information (causal, functional, generic, or nomological information) about the members of the concept's category of referents and their properties (e.g., Carey, 1985, 2009; Gopnik & Meltzoff, 1997). To make sense of that informational diversity, the present account builds on two hypotheses.

### 4.2.1 | Two pro-diversity hypotheses

The first hypothesis is the so-called 'heterogeneity' or 'non-uniformity hypothesis' (Machery, 2005, 2006, 2009; Weiskopf, 2009). It takes the three basic types of bodies of information, respectively correlated with exemplar-, prototype-, and theory-theories of concepts, to be different basic kinds of concepts, without any commonality. As such, these three basic kinds of concepts can be identified via a strong twofold individuation criterion, which is that, in order to belong to the same concept, two elements of information have to be: Connected, in the sense that the retrieval of one must facilitate access to the other (e.g., from water being

[17]Note that, in what follows, 'multiple realisation' does not much refer to the potential realisation of concepts *qua* functional kind by different neural structures as to their actual realisation by different cognitive structures.

[18]For the record, on the classical theory, concepts are made of definition-like structures that consist of separately necessary and jointly sufficient properties, which are taken to be both semantically analytic and epistemologically a priori and which supposedly spell out the individuation and application conditions of concepts.

$H_2O$ to its being transparent, odourless, etc.), and coordinated, in the sense that they must not yield equally authoritative conflicting judgements (e.g., 'tomatoes are fruits' and 'tomatoes are vegetables') (Machery, 2009; Machery & Seppälä, 2011). Thanks to its fine-grained and clear-cut distinction, such strong individuation criterion theoretically ensures a higher flexibility to the method of conceptual engineering on the diversity of our conceptual devices, as the standards and measures to assess and improve the cognitive efficacy of concepts will then be calibrated in a way that varies according to the kind of concept that is presently targeted.[19]

The second hypothesis may be dubbed the 'compatibility hypothesis.' It consists in taking the exemplar-, the prototype-, and the theory-theories of concepts as providing complementary accounts of the various ways in which the cognitive processes that underlie the exercise of our concept-involving higher cognitive competences work (Machery, 2005, 2006, 2009; Piccinini & Scott, 2006; Weiskopf, 2009), instead of seeing them as vying with one another to be the single correct paradigm for theorising about one supposedly unique and homogeneous/uniform kind of concepts. The empirical argument for the complementarity of the three main paradigms is simply that none of them alone can account for the different categorisation strategies that we use in our higher cognitive activities (see also Bloch-Mullins, 2018; Goldstone et al., 2018; Strohminger & Moore, 2010). Consequently, the corresponding theoretical entities of each of these paradigms (exemplars, prototypes, and theories) can in turn be seen as representing three basic types of bodies of information that may be retrieved by default when one exercises a concept-involving cognitive competence in the execution of some higher cognitive task. Thanks to its resulting all-inclusive approach to informational diversity, such conceptual pluralism ensures a wider scope to the method of conceptual engineering on the diversity of our conceptual devices, as nothing of our concept-involving cognition will fall out of its remit. These two hypotheses lead to a conceptual pluralism that is both radical and absolute.

## 4.2.2 | Radical and absolute conceptual pluralism

The form of conceptual pluralism that results from our two 'pro-diversity hypotheses' is both radical and absolute, since it maintains that, for any referent, for any competence, and according to different processes, we may retrieve and use any of the three basic kinds of concepts—exemplars, prototypes, and theories. For instance, we may have an exemplar, a prototype, and a theory concept of dog, all (loosely) co-referential to the class of dogs (then overdetermined), and all possibly retrieved and used in the categorisation, the deduction, the induction, etc. processes that underlie the exercise of a corresponding competence (viz., categorizing, etc., respectively) in the execution of a given higher cognitive task, such as judgment-making, and so on, about dogs. On this picture, conceptual engineering will then consist in a differentiated amelioration of the default bodies of information that concepts are, depending on what their basic kind is, the processes that these basic kinds play a role in, and the competence that these processes underlie in the execution of a given cognitive task. Besides fitting again very well with the cognitive functionalism that is one of the hallmarks of conceptual engineering, this characterisation of concepts as multiply realised in a radical and absolute fashion thereby theoretically secures and justifies the broad-spectrum ambitions of its research program when these are cashed out in terms of their scope and flexibility on our conceptual devices in our whole cognitive life.

---

[19]An alternative option for dealing with all the basic types of bodies of information is that offered by hybrid theories of concepts on which exemplars, prototypes, and theories are the composing parts of concepts then conceived as single, highly structured entities of an overall homogeneous/uniform kind (e.g., Gelman, 2004; Keil, 1989; Keil et al., 1998); however, as a result of their non-compliance with the connection-coordination individuation criterion, hybrid models presumably cannot secure the higher flexibility that is required for the method of conceptual engineering to precisely operate on the diversity of our conceptual devices, in a calibrated way (cf. Vicente & Martínez Manrique, 2016).

## 4.3 | Upshot

This section investigated the concept of concept at work in the psychological framework with a view to improving its conduciveness to the purposes of conceptual engineering guided by two main desiderata—namely, to fit well with the cognitive functionalism of conceptual engineering and to further support its broad-spectrum ambitions. These two desiderata have been satisfied by the development in two steps of a psychological characterisation of concepts as 'multiply realiz[ed] functional kinds' (cf. Lalumera, 2010). With this concept of concept at hand, the upshot for conceptual engineering as the method for the cognitive amelioration of our conceptual devices is that it will then consist in redesigning the mind's patterns of inferences, possibly at each step of information processing (that is, from retrieval to use) and at any level of abstraction for the informational content (that is, for any basic kinds of concepts). Each basic kind of concepts will thus be assessed and improved in accordance with how well it enables the exercise of a given higher cognitive competence—in other words, for its 'cognitive efficacy'—in a task-appropriate and context-dependent way. Because it thereby theoretically secures and justifies the strongest impact, the widest scope, and the highest flexibility for the method of conceptual engineering on our conceptual devices in our whole cognitive life, I contend that this psychological concept of concept as multiply realised functional kind is the most conducive to the purposes of conceptual engineering, on its broad-spectrum version.

## 5 | CONCLUSION

This paper has investigated the theoretical foundations of conceptual engineering, characterised as the method to assess and improve our representational devices, fitted with broad-spectrum ambitions in terms of scope, impact, and flexibility. In order to ensure the actionability of conceptual engineering as a broad-spectrum method, the main problem this paper has addressed was that of its subject matter, namely: What is the proper object domain of conceptual engineering? This problem in turn was subdivided into two successive issues: First, how to construe the representational devices that conceptual engineering target? Or more explicitly, should conceptual engineering target concepts? Second, if conceptual engineering is to target concepts, how to theorise concepts, then, for the purposes of conceptual engineering? In addressing these foundational issues stepwise, I came to successively defend the three following theses, narrowing down thereby the target domain of conceptual engineering:

1. The representational devices that conceptual engineering targets should be construed as concepts, on pain, otherwise, of pragmatic inconsistencies;
2. The concepts that conceptual engineering targets should be psychologically theorised, on pain, otherwise, of limiting the broad-spectrum ambitions for its method;
3. The psychological concepts that conceptual engineering targets should be theorised as multiply realised functional kinds in order to further foster their conduciveness to achieving the purposes of conceptual engineering, on its broad-spectrum version.

All in all, by establishing that conceptual engineering *should* target concepts, psychologically theorised, as multiply realised functional kinds, in order to theoretically secure and justify a wide scope, a high flexibility, and a strong impact for its method on our conceptual devices in our cognitive life, I claim to have fully substantiated a broad-spectrum version of conceptual engineering. May it now live up to its self-proclaimed transformational power, so that eventually: 'Through conceptual engineering, we can take some control over what we can think, say, do, and be' (Scharp, 2020, p. 398).

## ACKNOWLEDGEMENTS

## ORCID

*Manuel Gustavo Isaac* iD https://orcid.org/0000-0002-5479-5027

## REFERENCES

Andow, J. (2020). Fully experimental conceptual engineering. *Inquiry*, Online First, 1–27. https://doi.org/10.1080/00201 74X.2020.1850339

Bloch-Mullins, C. L. (2018). Bridging the gap between similarity and causality: An integrated approach to concepts. *The British Journal for the Philosophy of Science*, *69*(3), 605–632. https://doi.org/10.1093/bjps/axw039

Brun, G. (2016). Explication as a method of conceptual re-engineering. *Erkenntnis*, *81*(6), 1211–1241. https://doi.org/10.1007/s10670-015-9791-5

Burge, T. (1993). Concepts, definitions, and meanings. *Metaphilosophy*, *24*(4), 309–325. https://doi.org/10.1111/j.1467-9973.1993.tb00198.x

Burgess, A. (2020). Never say 'never say "never"'. Chap. 5. In A. Burgess, H. Cappelen, & D. Plunkett (Eds.), *Conceptual engineering and conceptual ethics* (pp. 125–131). Oxford University Press.

Burgess, A., & Plunkett, D. (2013). Conceptual ethics I. *Philosophy Compass*, *8*(12), 1091–1101. https://doi.org/10.1111/phc3.12086

Burgess, A., & Plunkett, D. (2020). On the relation between conceptual engineering and conceptual ethics. *Ratio*, *33*(4), 281–294. http://doi.org/10.1111/rati.12265

Cantalamessa, E. A. (2021). Disability studies, conceptual engineering, and conceptual activism. *Inquiry*, *64*(1–2), 46–75. http://doi.org/10.1080/0020174x.2019.1658630

Cappelen, H. (2018). *Fixing language: Conceptual engineering and the limits of revision*. Oxford University Press.

Cappelen, H. (2020). Conceptual engineering: The Master Argument. Chap. 7. In A. Burgess, H. Cappelen, & D. Plunkett (Eds.), *Conceptual engineering and conceptual ethics* (pp. 132–151). Oxford University Press.

Cappelen, H., & Plunkett, D. (2020). A guided tour of conceptual engineering and conceptual ethics. Chap. 1. In A. Burgess, H. Cappelen, & D. Plunkett (Eds.), *Conceptual engineering and conceptual ethics* (pp. 1–34). Oxford University Press.

Carey, S. (1985). *Conceptual change in childhood*. MIT Press.

Carey, S. (2009). *The origin of concepts*. Oxford University Press.

Chalmers, D. J. (2011). Verbal disputes. *Philosophical Review*, *120*(4), 515–566. https://doi.org/10.1215/00318108-1334478

Chalmers, D. J. (2020). What is conceptual engineering and what should it be? *Inquiry*, Online First, 1–18. https://doi.org/10.1080/0020174X.2020.1817141

Deutsch, M. E. (2020a). Speaker's reference, stipulation, and a dilemma for conceptual engineers. *Philosophical Studies*, *177*(12), 3935–3957. http://doi.org/10.1007/s11098-020-01416-z

Deutsch, M. E. (2020b). Trivializing conceptual engineering. *Inquiry*, Online First, 1–16. https://doi.org/10.1080/00201 74X.2020.1853343

Deutsch, M. E. (2021). Still the same dilemma for conceptual engineers: Reply to Koch. *Philosophical Studies*, Online First, 1–12. https://doi.org/10.1007/s11098-020-01546-4

Dutilh Novaes, C. (2016). Conceptual genealogy for analytic philosophy. In J. A. Bell, A. Cutrofello, & P. M. Livingston (Eds.), *Beyond the analytic-continental divide: Pluralist philosophy in the twenty-first century* (pp. 74–108). Routledge.

Dutilh Novaes, C. (2020). Carnap meets Foucault: Conceptual engineering and genealogical investigations. *Inquiry*, 1–27. http://doi.org/10.1080/0020174x.2020.1860122

Eklund, M. (2021). Conceptual engineering in philosophy. Chap. 1. In J. Koo & R. Sterken (Eds.), *Routledge handbook of social and political philosophy of language* (pp. 15–30). Routledge.

Fischer, E. (2020). Conceptual control: On the feasibility of conceptual engineering. *Inquiry*, 1–29. http://doi.org/10.1080/0020174x.2020.1773309

Fodor, J. A. (1975). *The language of thought*. Harvard University Press.

Gelman, R. (2004). Cognitive development. Chap. 12. In H. Pashler & D. L. Medin (Eds.), *Stevens' handbook of experimental psychology memory and cognitive processes* (Vol. *2*, pp. 533–560). Wiley.

Goldstone, R. L., Kersten, A., & Carvalho, P. F. (2018). Categorization and concepts. Chap. 8. In J. T. Wixted & S. L. Thompson-Schill (Eds.), *Stevens' handbook of experimental psychology and cognitive neuroscience. language and thought* (Vol. *3*, pp. 275–317). Wiley.

Gopnik, A., & Meltzoff, A. N. (1997). *Words, thoughts, and theories*. MIT Press.

Greenough, P. (2020). Neutralism and conceptual engineering. Chap. 11. In A. Burgess, H. Cappelen, & D. Plunkett (Eds.), *Conceptual engineering and conceptual ethics* (pp. 205–229). Oxford University Press.

Haslanger, S. A. (2020a). Going on, not in the same way. Chap. 12. In A. Burgess, H. Cappelen, & D. Plunkett (Eds.), *Conceptual engineering and conceptual ethics* (pp. 230–260). Oxford University Press.

Haslanger, S. A. (2020b). How not to change the subject. Chap. 12. In T. Marques & Å. Wikforss (Eds.), *Shifting concepts: The philosophy and psychology of conceptual variability* (pp. 235–259). Oxford University Press.

Isaac, M. G. (2020). How to conceptually engineering conceptual engineering? *Inquiry*, Online First, 1–25. https://doi.org/10.1080/0020174X.2020.1719881

Johnston, M., & Leslie, S.-J. (2012). Concepts, generics, and the Canberra Plan. *Philosophical Perspectives*, *26*(1), 113–171.

Jorem, S. (2021). Conceptual engineering and the implementation problem. *Inquiry*, *64*(1–2), 186–211. http://doi.org/10.1080/0020174x.2020.1809514

Keil, F. C. (1989). *Concepts, kinds, and cognitive development*. MIT Press.

Keil, F. C., Carter Smith, W., Simons, D. J., & Levin, D. T. (1998). Two dogmas of conceptual empiricism: Implications for hybrid models of the structure of knowledge. *Cognition*, *65*(2), 103–135. https://doi.org/10.1016/S0010-0277(97)00041-3

Koch, S. (2020). Engineering what? On concepts in conceptual engineering. *Synthese*, Online First, 1–21. http://doi.org/10.1007/s11229-020-02868-w

Koch, S. (2021). There is no dilemma for conceptual engineering. Reply to Max Deutsch. *Philosophical Studies*, *178*(7), 2279–2291. http://doi.org/10.1007/s11098-020-01546-4

Lalumera, E. (2010). Concepts are a functional kind. *Behavioral and Brain Sciences*, *33*(2–3), 217–218. https://doi.org/10.1017/S0140525X10000403

Löhr, G. (2020). Concepts and categorization: Do philosophers and psychologists theorize about different things? *Synthese*, *197*(5), 2171–2191. http://doi.org/10.1007/s11229-018-1798-4

Machery, É. (2005). Concepts are not natural kinds. *Philosophy of Science*, *72*(3), 444–467.

Machery, É. (2006). How to split concepts: A reply to Piccinini and Scott. *Philosophy of Science*, *73*(4), 410–418. https://doi.org/10.1086/516812

Machery, É. (2009). *Doing without concepts*. Oxford University Press.

Machery, É. (2010). Précis of *Doing Without Concepts*. With open peer commentary and author's response. *Brain and Behavioral Science*, *33*(2), 195–244.

Machery, É. (2015). By default. In S. Laurence & E. Margolis (Eds.), *The conceptual mind: New directions in the study of concepts* (pp. 567–588). MIT Press.

Machery, É. (2017). Conceptual analysis rebooted. Chap. 7. In *Philosophy within its proper bounds* (pp. 208–244). Oxford University Press.

Machery, É., & Seppälä, S. (2011). Against hybrid theories of concepts. *Anthropology and Philosophy*, *10*, 99–126.

Margolis, E., & Laurence, S. (1999). Concepts and cognitive science. In S. Laurence & E. Margolis (Eds.), *Concepts: Core readings* (pp. 3–81). MIT Press.

Margolis, E., & Laurence, S. (2007). The ontology of concepts-abstract objects or mental representations? *Noûs*, *41*(4), 561–593. https://doi.org/10.1111/j.1468-0068.2007.00663.x

Margolis, E., & Laurence, S. (2010). Concepts and theoretical unification. *Behavioral and Brain Sciences*, *33*(2–3), 219–220. https://doi.org/10.1017/S0140525X10000427

Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, *85*, 207–238. https://doi.org/10.1037/0033-295X.85.3.207

Nado, J. (2021). Conceptual engineering, truth, and efficacy. *Synthese*, *198*(S7), 1507–1527. http://doi.org/10.1007/s11229-019-02096-x

Peacocke, C. (1992). *A study of concepts*. MIT Press.

Piccinini, G., & Scott, S. (2006). Splitting concepts. *Philosophy of Science*, *73*(4), 390–409. https://doi.org/10.1086/516806

Pinder, M. (2019). Conceptual engineering, metasemantic externalism and speaker-meaning. *Mind*, Online First, 1–23. https://doi.org/10.1093/mind/fzz069

Plunkett, D. (2015). Which concepts should we use? Metalinguistic negotiations and the methodology of philosophy. *Inquiry*, *58*(7–8), 828–874. https://doi.org/10.1080/0020174X.2015.1080184

Plunkett, D. (2016). Conceptual history, conceptual ethics, and the aims of inquiry: A framework for thinking about the relevance of the history/genealogy of concepts to normative inquiry. *Ergo*, *3*(2), 27–62. https://doi.org/10.3998/ergo.12405314.0003.002

Prinzing, M. (2018). The revisionist's rubric: Conceptual engineering and the discontinuity objection. *Inquiry*, *61*(8), 854–880. https://doi.org/10.1080/0020174X.2017.1385522

Rey, G. (1983). Concepts and stereotypes. *Cognition*, *15*(1–3), 237–262. https://doi.org/10.1016/0010-0277(83)90044-6

Rey, G. (1985). Concepts and conceptions: A reply to Smith, Medin and Rips. *Cognition*, *19*(3), 297–303. https://doi.org/10.1016/0010-0277(85)90037-X

Richard, M. (2020). The A-project and the B-project. Chap. 17. In A. Burgess, H. Cappelen, & D. Plunkett (Eds.), *Conceptual engineering and conceptual ethics* (pp. 358–378). Oxford University Press.

Rosch, E. (1973). Natural categories. *Cognitive Psychology*, *4*(3), 328–350. https://doi.org/10.1016/0010-0285(73)90017-0

Rosch, E. (1975). Cognitive representations of semantic categories. *Journal of Experimental Psychology: General*, *104*(3), 192–233. https://doi.org/10.1037/0096-3445.104.3.192

Rosch, E., & Mervis, C. B. (1975). Family resemblance: Studies in the internal structure of categories. *Cognitive Psychology*, *7*(4), 573–605.

Sawyer, S. (2018). The importance of concepts. *Proceedings of the Aristotelian Society*, *116*(2), 127–147.

Sawyer, S. (2020a). Talk and thought. Chap. 18. In A. Burgess, H. Cappelen, & D. Plunkett (Eds.), *Conceptual engineering and conceptual ethics* (pp. 379–395). Oxford University Press.

Sawyer, S. (2020b). Truth and objectivity in conceptual engineering. *Inquiry*, *63*(9–10), 1001–1022. https://doi.org/10.1080/0020174X.2020.1805708

Scharp, K. (2013). *Replacing truth*. Oxford University Press.

Scharp, K. (2020). Philosophy as the study of inconsistent concepts. Chap. 19. In A. Burgess, H. Cappelen, & D. Plunkett (Eds.), *Conceptual engineering and conceptual ethics* (pp. 396–416). Oxford University Press.

Simion, M. (2018). The 'should' in conceptual engineering. *Inquiry*, *61*(8), 914–928. https://doi.org/10.1080/0020174X.2017.1392894

Smith, E. E., & Medin, D. L. (1981). *Categories and concepts*. Harvard University Press.

Sterken, R. K. (2020). Linguistic intervention and transformative communicative disruptions. Chap. 20. In A. Burgess, H. Cappelen, & D. Plunkett (Eds.), *Conceptual engineering and conceptual ethics* (pp. 417–434). Oxford University Press.

Strohminger, N., & Moore, B. W. (2010). Banishing the thought. *Behavioral and Brain Sciences*, *33*(2–3), 225–226. https://doi.org/10.1017/S0140525X10000476

Tanswell, F. S. (2018). Conceptual engineering for mathematical concepts. *Inquiry*, *61*(8), 881–913. https://doi.org/10.1080/0020174X.2017.1385526

Thomasson, A. L. (2020). A pragmatic method for conceptual ethics. Chap. 21. In A. Burgess, H. Cappelen, & D. Plunkett (Eds.), *Conceptual engineering and conceptual ethics* (pp. 435–458). Oxford University Press.

Vicente, A., & Martínez Manrique, F. (2016). The big concepts paper: A defence of hybridism. *The British Journal for the Philosophy of Science*, *67*(1), 59–88. https://doi.org/10.1093/bjps/axu022

Weiskopf, D. A. (2009). The plurality of concepts. *Synthese*, *169*(1), 145–173. https://doi.org/10.1007/s11229-008-9340-8