Mechanisms of skilful interaction: Sensorimotor enactivism & mechanistic explanation[1]

Jonny Lee and Becky Millar

Abstract

The mechanistic model depicts scientific explanations as involving the discovery of multi-level, organised components that constitute a target phenomenon. Meanwhile, sensorimotor enactivism purports to offer a scientifically informed account of perceptual experience as a skill-laden interactive relationship, constitutively involving both perceiver and world, rather than as an agent-bound representation of the world. Insofar as sensorimotor enactivism identifies an empirically tractable phenomenon—skilful agent-world interaction—and mechanistic explanation establishes the subpersonal components of this phenomenon, the two approaches allow for a fruitful division of labour in investigating perceptual experience. On closer inspection, however, two challenges arise. First, the 'representation challenge' arises because promising attempts to set out implementational details of our sensorimotor interaction with the world implicate cognitive representations, creating tension with sensorimotor enactivism's non-representational commitments. Second, the 'reconstitution challenge' arises when mechanistic explanation not only uncovers the components of some established phenomenon but plays a role in 'reconstituting' this phenomenon. This means that, through investigating mechanisms, perceptual experience may be reconceived such that its constituents are wholly organism-bound. We explore both challenges to the compatibility of mechanism and sensorimotor enactivism and examine possible solutions. The result is a clearer understanding of the tensions and opportunities for learning between the frameworks.

Key Words: sensorimotor enactivism, mechanism, mechanistic explanation, perceptual consciousness, perception, representation

[1] We are grateful for helpful comments from Louise Richardson, Joe Dewhurst, and Dan Calder on earlier drafts of this paper. We also thank several anonymous reviewers who helped us to improve this manuscript.

1. Introduction

The mechanistic model is an account of scientific explanation (e.g. Machamer et al., 2000; Glennan, 2002, 2017; Bechtel & Abrahamsen, 2005). Applied to cognitive science, the mechanistic model (or simply 'mechanism') states that cognitive phenomena, broadly construed, are explained by their underlying mechanisms. Such phenomena include perception. The mechanistic model is pertinent to any theory in cognitive science to the extent that it is now the dominant account of explanation.

Sensorimotor enactivism (SME), meanwhile, characterises perceptual experience as a kind of interactive *relationship* between perceiver and world, involving an implicit understanding of sensorimotor contingencies (SMCs): law-like relations holding between sensory information and bodily activity. The approach can therefore be construed as a form of *relationalism* (Ward, 2012; 2022)—an approach contrasted with representational accounts. Relationalism is a metaphysical thesis, according to which, perceptual experience is most fundamentally a *relationship with* the world rather than a *representation of* the world. SME also takes itself to be an empirically informed research programme, which prompts experimental work into phenomena such as change blindness (O'Regan, Rensink, & Clark, 1999; Rensink, O'Regan, & Clark, 2000) and colour perception (Philipona & O'Regan, 2006).

If SME is congruous with scientific explanation, and mechanism captures the nature of such explanations, one would expect compatibility between SME and mechanism.[2] On the face of it, this seems to hold: SME articulates the phenomenon to be explained (roughly, what perceptual experience is) and so provides constraints on the possible mechanisms for perception; mechanistic explanation provides implementational details for this phenomenon, chiefly, how organised components and their operations allow for skilful worldly interaction, involving mastery over sensorimotor contingencies. On closer inspection, however, there are two challenges to the compatibility between SME and the mechanistic model with regard to explaining perceptual experience.

First, promising accounts have been offered which begin to sketch possible mechanistic implementations of SME (e.g., Seth, 2014), but such implementations appear to implicate representations. Tension arises here with the supposedly non-representational nature of SME. This is the 'representation challenge'. Second, the mechanistic model indicates the possibility that our understanding of how perception is constituted may alter over the course of scientific investigation, such that the constituents of perception turn out to be wholly 'internal', that is, within the confines of the perceiving organism. Tension may arise here with SME's metaphysical commitment to relationalism, according to which conscious perceptual experience constitutively involves both subject and the external world. This is the 'reconstitution

[2] Of course, one may resist the dominance of the mechanistic model, for instance, by holding dynamical explanations to be the best way of understanding cognitive phenomena (e.g., Chemero & Silberstein, 2008). Even if dynamical explanations offer an alternative to mechanistic explanations, the relationship between mechanism and SME is important to explore given the ascendency of the former in the contemporary literature. Moreover, many claim that dynamical explanations are a part of or otherwise collapse into mechanistic explanations (for discussion, see Bechtel & Abrahamsen, 2010; Zednik, 2011; Kaplan, 2018; Lee, 2023). If true, dynamism would not offer an escape from the mechanistic model.

challenge'. This paper examines these two challenges and argues that both may be overcome. To this extent, mechanism and SME are compatible.

The paper proceeds as follows. §2 introduces the mechanistic model and shows how it has been used to substantiate the notion of cognitive representation. §3 introduces SME, emphasising its commitment to relationalism about perceptual experience, and establishes the apparently innocuous sense in which it is compatible with mechanism. §4 reveals the creeping tension between SME and mechanistic attempts to implement mastery over sensorimotor contingencies given the latter's appeal to representational mechanisms (the 'representation challenge'). We propose three solutions: (i) abandoning the relationalist assumptions of SME, (ii) denying mechanistic models really are committed to representations, and (iii) maintaining that perception can be both relationalist and involve subpersonal representation. §5 turns to the possibility that mechanistic explanation may arrive at a different understanding of what constitutes perceptual experience (the 'reconstitution challenge'). In response, we underscore that such tension is contingent on the results of scientific findings. We also take the opportunity to highlight pitfalls in discussing the constituents of perception. In particular, we caution against conflating distinct notions of constitution—a recognition of which may allow us to preserve both SME and 'internalist' conceptions of mechanistic composition.

2.1 Mechanism

The mechanistic model has emerged in recent years as the dominant philosophical account of how effective explanations do and should advance in the cognitive, brain and perceptual sciences (hereafter, 'cognitive science'). According to the mechanistic model, cognitive science explains by offering mechanistic explanations. Such explanations proceed by identifying underlying organised causal structures—mechanisms and the components which comprise them—that are responsible for a phenomenon under investigation (e.g., Machamer et al., 2000; Glennan, 2002, 2017; Bechtel & Abrahamsen, 2005).

Bechtel & Abrahamsen provide a relatively uncontroversial characterisation of mechanisms, writing:

> A mechanism is a structure performing a function in virtue of its component parts, component operations, and their organization. The orchestrated functioning of the mechanism is responsible for one or more phenomena. (2005, p. 423)

In other words, a mechanism consists of component parts and processes, organised so as to realise a phenomenon. The organisation of components (their relations to one another) and their operations (the activities they perform) are thus central to understanding how a mechanism realises a target phenomenon.

Components are organised by their spatial and temporal properties. Investigating the location, size, and orientation of components (spatial properties), in conjunction with the order, rates and duration of their activities (temporal properties) explains how a phenomenon is realised. Components and their activities are typically structured into

nested levels of organisation. For example, the components comprising the heart include valves, which themselves are composed, among other things, of interstitial cells which are composed of various organelles and so on. Hence, components are often mechanisms themselves, though organised components and their activities are always at a lower level of organisation than the phenomenon they realise.

It is often stressed that there are no mechanisms 'as such' but rather mechanisms *for* phenomena (Glennan, 2002). Mechanistic explanation thus critically involves identifying a phenomenon; different phenomena require different mechanisms and so different explanations (though some mechanisms may be responsible for more than one phenomenon; for discussion, see Herschbach & Bechtel, 2014). Sometimes, mechanistic inquiry involves re-identifying the phenomenon if our initial characterisation appears inadequate considering what we discover about the mechanism—as occurred, for instance, throughout the history of investigating the mechanisms of memory (Craver, 2007; Bechtel, 2008; see §6). Thus, whilst a mechanism is individuated by the phenomenon it explains, how we understand the nature of the phenomenon we are explaining may evolve. As we explore below, one plausible way to understand the relationship between SME and mechanism is in terms of the former articulating the target phenomenon pertaining to perceptual experience in a manner amenable to mechanistic explanation; mechanistic investigation is compatible with SME to the extent that the phenomenon of perceptual experience is identified in terms of skilful interaction between perceiver and world.

Mechanistic explanations typically explain a phenomenon by uncovering its spatiotemporal constituents. It is thus worth clarifying the mechanistic conception of constitution. The mechanistic framework sets out a technical notion of constitution that states *x* is partially constitutive of a phenomenon *P* if *x* is an internal part of *P* at some point during *P*'s existence—phenomena have durations and parts/operations need only exist for as long as they are functionally operative. For example, the 'main wheel' is a constituent of an analogue clock because it is an internal part of the mechanism (the whole clock) relative to, say, the phenomenon of moving hands, indicating the time of day.

Under at least one interpretation, mechanistic explanation subsumes three types of explanation: 'etiological', 'contextual' and 'constitutive' (cf., Craver, 2007; Glennan 2009; Piccinini, 2020). These can also be framed in relational terms. Two or more items can bear one of three types of relation: etiological, contextual, and constitutive. Etiological explanation concerns the 'external' and historical causes of the phenomenon (e.g., assembling a clock, placing it on an appropriate surface, winding it up etc.). Contextual explanation concerns what role the phenomenon plays in a wider phenomenon (e.g., what effect a working clock has on a wider timekeeping system). Constitutive explanation concerns the causal structure that belongs to the phenomenon itself, that is, the internal spatio-temporal parts that the phenomenon consists of 'from the inside out' (e.g., the main wheel, escape wheel, hands, case, barrel etc.). This taxonomy of explanations and corresponding relations makes clear that constitution does not exhaust the factors that are causally necessary for a phenomenon; exposure to sunlight, for instance, is a necessary external cause of seed growth (but does not *constitute* the growing seed) (Piccinini, 2020).

A pertinent feature of mechanistic approaches to cognition is the absence of a priori constraints on the spatiotemporal location of a cognitive capacity's constituents, so it does not automatically preclude accounts according to which perceptual consciousness extends into the environment. For instance, a perceptual mechanism extends into the body and world if the constitutive parts of a cognitive phenomenon are spatially located in the body and world.[3] Thus, the distribution of a phenomenon's constitutive parts is a contingent fact, to be discovered empirically. In reality, however, determining where the boundaries of complex, dynamic phenomena begin and end is not straightforward. In partially addressing this difficulty, some mechanists have proposed an account of constitutive relations in terms of 'mutual manipulability' (Craver, 2007). In Craver's original formulation, a component *x* of *S* is constitutively relevant for *S* if and only if:

> (i) *x* is part of *S*; (ii) in the conditions relevant to the request for explanation there is some change to *X's φ*-ing that changes *S*'s *ψ*-ing; and (iii) in the conditions relevant to the request for explanation there is some change to *S*'s *ψ*-ing that changes *X*'s *φ*-ing. (Craver, 2007, p. 153)

Thus, intervening on a component will affect the behaviour of the mechanism as a whole, whilst intervening on the behaviour of the mechanism as a whole will affect the component. We will return to this notion below.[4]

2.2 Representational mechanisms

The role and value of 'representation' for explaining cognition is a hotly contested issue (for a good starting point, see Ramsey, 2007). Attempts have recently been made to clarify the notion by situating it within the mechanistic framework of explanation. This has been viewed as an opportunity to refine this often ambiguous and highly contested notion and clarify the empirically accountable conditions for its legitimate ascription (e.g., Gładziejewski, 2015; Lee, 2019, 2021; Bielecka & Miłkowski, 2020; Piccinini, 2020). In short, according to a mechanistic perspective, legitimate representation ascriptions in cognitive science refer to a type of mechanism individuated by a functional profile. This sufficiently resembles a class of ordinary representations such that 'representating' appropriately describes its operation. In turn, this role is typically understood in terms of a representational vehicle 'standing-in for' something (e.g.,

[3] As Piccinini (2020) puts it: 'Obviously, ordinary cognition is embodied—that is, dynamically coupled with a body. Equally obviously, ordinary cognition is embedded—that is, dynamically coupled with an environment. Perhaps ordinary cognition is not just embodied and embedded but extends into the environment as well—that is, parts of an agent's environment are also part of that agent's cognitive processes (e.g., Wilson 1994; Clark and Chalmers 1998)' (p. 248).

[4] Of relevance, mutual manipulability has been deployed by Kaplan (2012) to inform the extended mind debate. Kaplan's claim is that whether part of the agent's environment is a constitutive part of their cognition, can be understood in terms of whether it's part of a cognitive mechanism, i.e., the relationship of mutual manipulability between components of the cognitive system demarcates its boundaries. Note that mutual manipulability has not been an uncontroversial idea among mechanists and has spawned much discussion (e.g., Fazekas & Kertész, 2011; Leuridan, 2012; Krickel, 2018).

Ramsey, 2007). For instance, a cartographic map plays the functional role of representing, say, a mountain range, in virtue of standing-in for it.

The challenge then remains to identify how a cognitive mechanism could play the role of standing-in for something in a manner sufficiently analogous to ordinary representations. Here we can turn to the 'structural representation' account, or 'S-representation', for short (e.g., O'Brien & Opie, 2004; Gładziejewski, 2015; Gładziejewski & Miłkowski, 2017; Shea, 2018; Williams & Colling, 2018; Lee, 2019; Piccinini, 2020). The S-representation account stresses that a mechanism represents if it plays a map or model-like role for a system, that is, by instantiating a structural correspondence (isomorphism or homomorphism) with some task-relevant item (or 'target') which can be exploited to guide behaviour. This model-like mechanism, in turn, is sensitive to 'error', insofar as the containing system is sensitive to functionally significant mismatches in correspondence between the mechanism and target and accordingly adjusts the mechanism's functionally relevant structure (and thus the degree of correspondence). Example mechanisms that appear to meet the requirements for S-representation include 'cognitive maps'—neural mechanisms for navigation that operate (in part) by exploiting structural similarities between features of the mechanism (primarily located in the hippocampus) and the organism's environment, for instance, the nomic correspondence between the strength of the connections between the cells represents and distances between points in the environment (Tolman, 1948; O'Keefe & Dostrovsky, 1971; O'Keefe, 1976).

Crucial for grounding the distinctly representational quality of explanations that appeal to S-representations is the fact that the degree of correspondence between the mechanism and some behaviourally relevant target (e.g., cognitive map and environment), is crucial for explaining the degree of behavioural success (e.g., in locating some food source), in addition to the capacity of the system to update the mechanism following behavioural failure (e.g., Gładziejewski & Miłkowski, 2017).

In summary, the S-representation account is intended to specify a functional profile for a mechanism to qualify as representing, or in other words, as standing-in for something. If a mechanism meets the requirements for S-representation, the idea goes, it functions as a representation in a robust sense. As we shall see in §4, attempts have been made to undermine the representational status of even mechanisms meeting the requirements for (so-called) S-representation. This is pertinent to assessing the relationship between SME and mechanism because it transpires that some accounts of how to implement SME appear to implicate S-representations. First, however, we turn to the core commitments of SME and its potential compatibility with mechanistic explanations of perceptual experience.

3.1 Sensorimotor enactivism

SME claims that perceptual experience is to be understood in terms of skilful engagement with the world on the part of the whole animal. Such engagement is mediated by 'sensorimotor understanding': an implicit attunement to, or expectations about, *sensorimotor contingencies* (SMCs). SMCs are the systematic ways that our bodily activities induce sensory changes. They include, for example, the characteristic

ways that sensory stimulation on the retina changes as the eye rotates. This is predictable and is based upon 'the size of the eye movement, the spherical shape of the retina, and the nature of the ocular optics' (O'Regan & Noë, 2001, p. 941). According to this approach, it is through attunement to the ways that our bodies can induce sensory changes that we are able to successfully perceive a stable, non-gappy world of complete objects that goes beyond the immediate sensory information that we receive.

The sensorimotor approach is supported by the various ways in which perceptual experience seems to outstrip our current sensory contact with the world. For example, according to Noë (2004, pp. 59–65), SME has the advantage of solving the puzzle of 'perceptual presence'. This is the puzzle of how hidden aspects of objects and scenes—'that which, strictly speaking, I do not perceive' (ibid., p. 60)—can make their way into perceptual experience. We perceive a tomato as spherical, but how can this be when we do not have sensory contact with the back of it? SME offers a solution: the hidden parts of objects are present as accessible. Through our skilful sensorimotor interaction with the world, we have access to the whole, spherical tomato. Implicitly understanding how our own possible movements can provide sensory access to different parts of the environment enables an experience of richly detailed scenes and whole objects, even though in an important sense we do not have current perceptual awareness of all these details at once.

SME has been taken to be in some sense inimical to representations. For instance, it holds that we do not need neurally-realised models of the world's detail because the detail is already there in the world, ready to be accessed through skilful bodily activity. This idea is sometimes expressed through the claim that the world serves as an 'outside memory' (O'Regan, 1992); through bodily activity, we access the world's detail as and when required.[5] Rejecting the need for the construction of a richly detailed model of the environment puts the sensorimotor approach in opposition to orthodox cognitive science (e.g., views descending from Marr, 1982; see discussion in O'Regan & Noë, 2001, p. 940). Thus, at the very least, neural representations appear to play a more minimal role in SME than in orthodox accounts.

These observations about the nature of our access to environmental information lead to a second, closely related reason why SME might be seen as unfriendly to representations. Specifically, supporters of SME often take our skilful interactions with the world to put us in a direct perceptual relationship with worldly objects (e.g., Hurley, 1998; Noë, 2004; 2008; 2009; Beaton, 2016; Ward, 2012; 2022). Thus, although this connection isn't always explicitly spelt out, the approach is plausibly best construed as a *relational* account of conscious perceptual experience (see Ward, 2012; 2022 for detailed discussion). Relational approaches characterise veridical perceptual experience as most fundamentally some form of relationship *with* the world, rather than a representation *of* the world (e.g., Campbell, 2002; Martin, 2004; Brewer, 2011), and SME provides a story about how, through our implicit understanding of SMCs, such a relationship with the world is possible. So understood, veridical perceptual experience essentially involves both the perceiver and the worldly objects and properties to which they bear this skill-mediated relationship; it is therefore, as Ward (2012) puts it, a form

[5] C.f. Brooks' (1991) claim that the world is its own best model.

of ‘conscious externalism’. The (whole, spherical) tomato itself figures as a constituent of perceptual experience insofar as the perceiver is appropriately related to this object through their sensorimotor skills.

While SME provides a metaphysical account of the essential nature of perceptual experience, understood as a personal-level phenomenon, it is also intended to be scientifically informed and to guide cognitive science. For instance, it appears to be supported by, and has prompted work on, change and inattentional blindness (for discussion see O’Regan & Noë, 2001). The appearance of distractors such as ‘mud splats’ (O’Regan, Rensink & Clark, 1999), or the provision of an attention-intensive task (e.g., Simons & Chabris, 1999), can cause people not to notice large changes or unusual events within a scene. These phenomena may be surprising if it is assumed that perceptual experience is explained through a richly detailed model of the environment. They are rendered unpuzzling, however, by instead treating perception as involving skilful *access* to this detail, rather than a current representation of—or sensory contact with—such detail. Other empirically relevant phenomena that SME has prompted research into, and discussion of, include colour perception (e.g., Philipona & O’Regan, 2006), neural plasticity (e.g., Hurley & Noë, 2003), and the impact of visual reversing goggles upon perception (e.g., Degenaar, 2014). Attempts to treat SME as a scientific research programme or, at least, as continuous with science, raise questions about how it fits with the mechanistic model of scientific explanation.

3.2 SME’s compatibility with mechanism

The relationship between metaphysical and scientific accounts of perceptual experience is controversial (e.g., see Drayson, 2021; Fish, 2021; Noordhof, 2021 for recent discussion), but if SME is continuous with science, we would expect it to be compatible with scientific explanation. The mechanistic model provides an account of scientific explanation. Assuming the model is correct, we would expect SME to be compatible with mechanistic explanation. The remainder of this paper is dedicated to exploring the degree to which this is the case. In this section, we set out an innocuous sense in which SME and mechanistic explanation are compatible. In the following sections, we problematise this picture by raising, then countering, two issues that arise when integrating SME and mechanistic explanation.

SME and mechanism appear, ostensibly, to be compatible. This can be seen by noting that SME is principally concerned with providing an account of perceptual experience at the *personal level* (i.e., as a relationship holding between the subject and relevant parts of the world). This allows for a division of labour between SME and mechanistic explanation: SME identifies the personal-level phenomenon to be explained whilst mechanistic explanations show how this phenomenon is realised by organised causal parts and their operations. SME is arguably well-placed to specify a scientifically tractable phenomenon for mechanistic investigation because it purports to deflate alleged problems of perceptual consciousness, taking such experiences simply to be

episodes of skilful interaction with the environment.[6] O'Regan provides the illustrative example of the experience of a sponge's softness:

> Consider for example the "hard" question of how the softness of the sponge might be generated by brain mechanisms. From the point of view of the sensorimotor theory, this question does not make sense, because the softness lies in the fact that if I press the sponge, it squishes under my pressure. (2014, p. 25)

One need not appeal to 'qualia' (e.g., see O'Regan & Noë, 2001) since the experience of softness is nothing more than an interactive relationship with an object, mediated by an understanding of the sensory implications of one's own bodily movements. SME thereby offers a route for mechanistic investigation of perceptual consciousness—something that with certain approaches may appear mysterious, private, ineffable or more generally unsuitable for scientific investigation (c.f., Dennett's 1988 critique of qualia). With the sensorimotor approach, those conducting a mechanistic explanation of perceptual consciousness simply need to set out the mechanisms underlying the (empirically tractable) notion of a skilful interactive relationship between agent and world.

Drawing upon Ward (2012), we can see that due to this division of labour, even the most controversial aspects of SME (such as its relationalism and conscious externalism) are compatible with mechanistic explanation in a fairly trivial sense. Since SME specifies a personal-level phenomenon that extends between subject and world (i.e., an interactive relationship), if mechanistic explanation is tasked with setting out the implementational details of this phenomenon, it will follow that the mechanisms underlying this phenomenon will likewise extend beyond the boundaries of the body or brain. This is because, as Ward highlights, our understanding of what's going on at the personal and subpersonal levels constrain one another, so there is no reason to think that 'the relevant parts of the world should simply drop out when we shift our focus to the subpersonal' (p. 8). Thus, '[I]f we conceive experience as a dynamic relationship between subject and environment, the temptation to go internalist when describing experience at the subpersonal level simply will not arise' (p. 11).

Consider again O'Regan's sponge example. Tasks for the mechanist in setting out the subpersonal underpinnings of an episode of agent-sponge interaction will include, for example, uncovering the mechanisms involved in one's fingers moving and grasping an object, those that allow the sponge to compress and expand, and any neural mechanisms that underpin the agent's ability to successfully interact with the sponge. It follows from SME's relational personal-level account, that whatever unpins one's dynamic relationship with the world will have to extend into the environment.

To summarise, we have so far set out a straightforward sense in which SME and mechanism are compatible: SME clarifies the phenomenon of perceptual experience to which can then be investigated for underlying mechanisms. However, this relationship

[6] For example, versions of the approach have aimed to address the absolute explanatory gap or 'hard problem' of consciousness (O'Regan & Noë, 2001; O'Regan, 2014; Ward, 2022) or at the very least deflate the so-called 'comparative' explanatory gaps—i.e., the question of why experience is like *this* rather than like *that* (Hurley & Noë, 2003).

is less frictionless than first appears. In the following sections, we will first investigate an apparent tension between SME and the involvement of representational mechanisms in perception, before turning to possible disparities between sensorimotor enactivist and mechanist conceptions of where the constituents of perception lie.

4.1 The representation challenge

SME is often held to be opposed to representational views. In particular, it typically characterises perception in essentially relational rather than representational terms. This may be thought to constrain the kinds of subpersonal underpinnings of perception that are acceptable for the sensorimotor theorist. Since sensorimotor theorists usually focus on the personal level, they generally do not provide details about the mechanisms involved at the subpersonal level. However, promising accounts have been developed that do purport to flesh out these details for SME. However, problematically, they seemingly involve *representational* mechanisms that may be thought to 'contaminate' SME's essentially relational view of perception.

Consider the 'predictive processing' (PP) approach as an example. According to Seth (2014), PP can help illuminate the hitherto missing details of how mastery over SMCs is implemented in the brain (cf. Seth, 2015; Seth & Tsakiris 2018). PP is principally a theory of brain organisation according to which the primary function of neural activity is minimising errors in its endogenously produced anticipations of sensory stimulation. In brief, the brain operates in a constant cycle of generating and updating expectations of sensory input. To do so, it continually updates a model of the body and environment, which is used to generate predictions that are measured against sensory stimulation. Perception is achieved, at least in part, via evaluating top-down predictions about sensory stimuli against incoming sensory input, across multiple layers of a hierarchical processing architecture (for an overview of the neurophysiological evidence, see Walsh et al, 2020). In this way, 'perception involves the use of a unified body of acquired knowledge (a multi-level "generative model") to predict the incoming sensory barrage' (Clark, 2015, p. 5). This story, in turn, implicates particular mechanisms and acts as a research heuristic for discovering multilevel mechanisms (for discussion on PP and mechanistic explanation, see Gładziejewski, 2019; Harkness, 2015).

In combining aspects of SME with PP, Seth (2014) draws on behavioural and physiological evidence to suggest that perception involves 'counterfactually rich' generative models. In effect, the anticipatory machinery realised in the brain encodes both the likely causes of sensory input as well as the likely causes of sensory inputs predicted to occur given *possible* (though not necessarily realised) actions. The phenomenon of perceptual presence, in particular, depends on the 'counterfactual richness' of these generative models. Perceptual presence is absent or weakened when generative models fail to encode the likely causes of stimulation following action, hence is 'counterfactually poor'. The generative model is said to not only predict incoming sensory information but also how sensory information would change given the execution of actions available to the system.

Though he does not state so himself, Seth's (2014) account helps provide a mechanism sketch for the mastery of SMCs in accordance with the S-representation

account, outlined above.[7] Indeed, proponents of the S-representation approach to understanding the role and value of representation in cognitive science have appealed to PP as a case study. The basic idea is that for the brain to generate a prediction of the sensory signal it must instantiate a causal-probabilistic structure that maps onto the hidden worldly causes of stimuli, accomplished through a mechanism encoding a multi-level network of updatable priors (Gładziejewski, 2016).[8] As Vázquez (2020) summarises in her own discussion of Seth's account, generative models track structural features of the environment, thus: 'Generative models can be characterised as structural representations (S-representations) because […] they track the causal structure of what is represented in agent-dependent terms' (p. 668).

Accounts like Seth's raise the possibility that the mechanisms for mastery over sensorimotor contingencies involve representations.[9] This generates a prima facie puzzle: *if* an account of sensorimotor contingencies like Seth's is correct, how can we reconcile the relationalist (non-representational) commitments of SME with the involvement of representational mechanisms? If sensorimotor contingencies are understood as representational mechanisms, then it would seem SME is open to perceptual phenomena being explained in terms of representations. Thus, there are tensions between mechanistic explanation and SME.

Sensorimotor theorists often take it to be relatively unproblematic to provide an account that is non-representational at the personal level while still allowing some role for representations at the subpersonal level (e.g., O'Regan & Noë, 2001; Noë, 2004). For instance, Noë writes,

> No doubt perception depends on what takes place in the brain, and very likely there are internal representations in the brain (e.g., content-bearing internal states). What perception is, however, is not a process in the brain but a kind of skilful activity on the part of the animal as a whole (2004, p. 2)

Despite such ecumenical ambitions, however, one may worry that if perception of some object involves a robust subpersonal-level representation of that object, this must 'infect' the personal-level account of what it means to consciously experience it. If some subpersonal component of conscious experience is representing the world, then one might worry that conscious experience is representational too; the whole phenomenon cannot exclude this property of one of its parts. In short, the concern is that any reliance on a 'stand-in' for the perceptual object compromises the essentially relational rather than representational nature of perception. If subpersonal mechanisms

[7] For a recent argument that performing any complex control functions (as the nervous system does) requires S-representations, see Piccinini (2020).

[8] The role of representation for PP has been the subject of much debate, beyond Seth's particular account. For sample discussion, see Gładziejewski (2016); Dolega (2017); Sims & Pezzulo (2021); Rutar, Wiese & Kwisthout (2022).

[9] Of course, one may be unconvinced of Seth's account, and there are non-representationalist renditions of predictive processing (e.g., Kirchhoff & Robertson, 2018). Our point is simply that Seth provides an empirical account of sensorimotor contingencies—one that many may find compelling—which implicates representational mechanisms.

are map/model like, then agents are not directly related to the objects of perception (which the subpersonal mechanisms represent).

This concern can be seen more clearly through consideration of where exactly the distinction between relationalism and representationalism is situated. This distinction is sometimes presented not in terms of *whether* the perceiver is related to the world, but rather, in terms of the *type* of relation holding between perceiver and world. For instance, Locatelli and Wilson (2017) take the orthodox representationalist to concur with the relationalist about perception essentially involving a relation to the world, but contra the relationalist, this holds *by virtue of* representing the world to be some way. With such a distinction, then, relationalism holds that the perceiver's relation to the world is not an essentially representational one. But if there are robust representations at the subpersonal level, in what sense can this relation remain non-representational?

Such 'infection' worries may be amplified if one thinks that representations are required to do especially 'heavy lifting' for certain types of perceptual phenomena. For instance, Nanay argues that our best explanations of unconscious perception and multimodal integration rest upon representations. If we want to give an account of 'perception per se' (Nanay, 2014), we might wonder why a relational account is warranted at the personal level if representations are used to explain other important perceptual phenomena. Why not give a unified representational account of both?

Although—as we have detailed—sensorimotor theorists generally focus on perceptual consciousness, and so may not be unduly worried about accounting for non-conscious phenomena, accounts like Seth's appear to place a heavy explanatory burden upon representations even for *conscious* perception.Consider again Seth's account of *perceptual presence* (i.e., a personal-level phenomenon) in terms of a counterfactually-rich generative model. This account might be understood as one according to which our (personal-level) experiences of whole voluminous objects are a matter of neurally representing certain contents (i.e., content relating to the not currently present sides of object) such that these contents make their way into conscious experience. Seth's account, so construed, appears to conflict with SME's explanation of perceptual presence in terms of an agent bearing a relationship of skilful interaction to whole objects and their properties.

4.2 Solutions to the representation challenge.

We contend that there are at least three options in response to the representation challenge that facilitate the integration of SME and mechanism. These options are: (1) abandoning the relationalist commitments of SME; (2) embracing deflationism—the view that the S-representations implicated by accounts like Seth's are not, in fact, genuine representations; and (3) denying relationalism and representationalism, in the operative senses, are incompatible. Whilst our primary intent is to map these options, we take the third to be the most conciliatory solution.

4.2.1 Abandoning relationalism

One way to resolve the representation challenge is to abandon relationalism. This solution is simple: if we reconstrue SME so that it is not committed to relationalism, then there is no tension between the apparent existence of representational mechanisms and SME. This option involves reinterpreting SME so that it retains some of its core identity—namely, its commitment to perceptual experience involving mastery over SMCs—whilst dropping the claim that veridical perceptual experience is an interactive relationship between perceiver and world. Such an approach may take perceptual experience to be constituted purely by the internal representations involved in such a mastery of SMCs. Seth's marriage of SME and PP could be construed as exemplifying an internalist and non-relational version of SME, insofar as he emphasises that, 'the local material vehicles of perceptual experiences remain fully intracranial' (2014, p. 103).

There are, however, costs for the sensorimotor theorist in abandoning relationalism. First, understanding perceptual experience as an interactive relationship seems core to SME's attempts to defuse puzzles of consciousness. For instance, taking the experience of the sponge's softness to be nothing more than an episode of skilful interaction with the sponge—i.e., as something that spreads between brain, body and world—is the aspect of the account that is supposed to render the conscious experience unmysterious.[10] Second, and relatedly, the relational approach also seems central to attempts to characterise puzzling features of perceptual phenomenology, such as 'perceptual presence', whereby experience outstrips our sensory contact with the world (e.g., as in Noë, 2004). As Ward helpfully puts it:

> Though the perceptual information present in our sensory uptake from a particular time and place is impoverished in various ways, the enactivist holds that our understanding and expectations concerning the ways in which changes in this information march in step with changes in our relationship to the environment allow us to transcend these limitations and be related to the world, not just its impacts on our sensory receptors, in experience. (2012, p. 6)

With sensorimotor relationalism, it is one's interactive relationship with the world that does the explanatory work here. While there might be ways for a sensorimotor theorist to address these puzzles without recourse to relationalism, extant promising solutions provided by SME appear to be best understood in relationalist terms (see Ward, 2012).

As we shall discuss, the upshot of holding on to sensorimotor relationalism is that, assuming S-representations really do play a role in realising perceptual experience, one must either deny they really are representations (§4.2.2) or find some other method of reconciliation (§4.2.3).

[10] Ward (2022) offers a new framing of how sensorimotor relationalism can resolve puzzles of consciousness. He argues that sensorimotor theory enables an explanation of consciousness *tout court*, while relationalism allows for an explanation of why *particular* experiences are associated with *particular* phenomenal properties (i.e., because these properties are identified with particular mind-independent properties). However, for Ward, both sensorimotor and relational aspects are needed for a satisfactory and naturalistically tractable account of conscious experience.

4.2.2 Deflationism

A second way to resolve the tension between SME's commitment to relationalism and the possibility of representational mechanisms underlying perception is to deny that the latter amount to representations in a robust sense of the term, or similarly, that the so-called representational mechanism in question are only representations in a highly deflated sense. Such an approach may appeal to explicitly anti-representational versions of enactivism, such as 'radical enactivism' (Hutto & Myin, 2013).[11]

Some accounts of the subpersonal mechanisms involved in perception may use the term 'representation' even though the putative mechanisms under consideration do not function as a map, model, or other intermediary in perception. Much has been made in recent years of the fact that 'representation' has been applied loosely and with varying standards for attribution across cognitive science. Ramsey (2007), for instance, has argued that many of the entities labelled as representations do not function in a sufficiently representation-like manner (cf. Ramsey, 2017). For example, neural mechanisms which merely fire reliably in the presence of some stimuli cannot be said to 'stand-in for' anything and so may be better described as 'relays', 'detectors' or 'causal mediators' rather than representations (Gładziejewski & Miłkowski, 2017). Such mechanisms are only 'representations' in a highly deflated sense. Similarly, some critics of representation have argued that theories of representation (like teleosemantics) provide a theory of 'intentional directedness' without amounting to representational content, and cannot form the basis for truth or accuracy conditions (Hutto & Myin, 2013). Thus, appeal to so-called representational mechanisms at the subpersonal level, or mere use of the representational label, does not necessarily evidence representations, in a robust sense. If it transpires that subpersonal 'representations'—as deployed in theories like PP that aim to provide mechanistic details for SME—do not function as stand-ins in the robust sense intended by the S-representation account, then they appear to pose no risk to a relational conception of perceptual experience.

The S-representation account is, of course, intended to provide criteria for a mechanism to count as a representation in a non-deflationary sense. The general idea is that certain ordinary representations (maps, models, etc.) are characterised by a particular functional profile that can be mirrored by cognitive mechanisms.[12] Cognitive maps and other mechanisms play a map or model-like role by guiding behaviour via structural correspondence with task-relevant items, and can be said to bear the 'accuracy conditions' characteristic of representations insofar as their degree of correspondence with those features determines the success (or degree of success) of the behaviour in question. In other words, the S-representation account has been developed, in part, in response to worries that so-called cognitive representations may not function as genuine representations, and to show how a subpersonal mechanism

[11] However, according to Ward (2022) the relationalist construal of SME may be incompatible with Hutto & Myin's view on other grounds. Namely, they reject appeal to implicit sensorimotor understanding, which, for the sensorimotor relationalist, is what enables one's perceptual relation to the world.

[12] Ramsey (2007) specifically identifies the S-representation account as successfully articulating criteria for representation worthy of the name i.e., S-representations *do* count as robust representations.

may operate as a stand-in, in a robust sense (e.g., Ramsey, 2007). Nevertheless, some have insisted that the S-representation account fails in this regard, and that 'similarity-based' conditions do not really imply the kind of content that genuine representation requires.

A version of this view has been articulated by Segundo-Ortin & Hutto (2019) who argue that proponents of the S-representation account 'presuppose' but do not 'explain' the origin of content. They summarise the logic of the S-representation account as follows:

> The properties of a given S-representational vehicle, R, cause it to be structurally similar to some target state of affairs, T. Because R can mirror the structure of T more or less accurately, structural similarity entails accuracy conditions. Accuracy conditions are taken to entail content. Therefore, structural similarity is taken to entail content. Thus, S-representationalists conclude, the fact that R structurally mirrors T entails that R contentfully represents T. (*ibid,* S10)

They proceed to argue, however, that structural similarities are not really contentful because whilst *we* can make truth-evaluable claims regarding the structural correspondence between two items, it does not follow one item contentfully represents the other in the sense that the former can say something true or false (*ibid.* S13). Segundo-Ortin & Hutto thus do not deny the importance of structural correspondence for cognition, but deny it serves as the basis for representational content.

If S-representations are not really representations in a robust sense, because they cannot ground representational content, then it would seem the representational credentials of accounts like Seth's is undermined and so the tension between SME and mechanistic models of mastery over SMCs is eased. Proponents of the S-representation account remain convinced, however, that it provides robust, naturalistic grounds for attributing content. We cannot go into the weeds of this debate, for reasons of space. For our purposes, it is enough to acknowledge that defences typically hinge on accepting that structural similarity is insufficient for content, whilst emphasising that accuracy conditions are pertinent given the full suite of conditions that S-representations fulfil (for sample discussion, see Lee, 2019; Bielecka & Miłkowski, 2020; Piccinini, 2022). Briefly, S-representations imply content, and are thus 'robust', because of the way their structural correspondence with task-relevant items determines behavioural success—much like an ordinary map or model does—further supported by the containing system's sensitivity to 'error' i.e., inadequate structural correspondence with task relevant items—much like a mountaineer is sensitive to inadequate correspondences between their map and environment following failure to successfully navigate.

Though deflationism remains a live option, we note that it depends on the success of a deeper, adjacent debate about the credentials of the S-representation account. The proponent of SME may be independently convinced of the S-representation account or may otherwise not wish to depend on the outcome of this debate. We now turn to the possibility that relationalism (in the sense important to SME) and representationalism (in the sense important to mechanism) are compatible.

4.2.3 Reconciling subpersonal representations with relationalism

A third way to resolve tension between the relational interpretation of SME and the possibility of representational mechanisms underlying perception is to deny that subpersonal robust representations must 'infect' the personal level, even if they do important explanatory work. Such an approach requires there to be a robust sense in which neural mechanisms can act analogously to ordinary representations, without this implying *perceptual experience* has representational status. In other words, the existence of S-representations at the subpersonal level does not make it the case that perceptual experience itself is helpfully described as representing the world as being such-and-such a way. This would allow the perceptual relationship between subject and world to be non-representational in an important sense.

There is conceptual space for such an approach. Neural mechanisms playing a map or model-like role may partially underly mastery over SMCs, without this necessitating that experience itself represents the world or that the kind of relationship holding between subject and world is best construed in representational terms. This is possible, we suggest, because even if representational mechanisms are necessary for perceptual experience, this does not imply that perceptual experience is not constituted by interactive relations between agent and object. SME is compatible with representational mechanisms, so long as (a) perceptual experience is not *wholly* constituted by a representation, and (b) representations partially facilitate the interactive relationship that *is* wholly constitutive of perceptual experience. At least some plausible readings of accounts like PP, we suggest, are compatible with (a) and (b). Although, as mentioned above, Seth (2014) appears to endorse internalism about perceptual experience, this does not appear to be essential to PP. It's not clear that the generative model that PP postulates needs to be construed as alone constitutive of conscious experience. It might be understood as only offering an account of certain important elements of the subpersonal underpinnings of perceptual experience. For instance, plausibly it provides an account of how the brain enables mastery of SMCs, which for the sensorimotor relationalist is just one aspect of the interactive relationship constitutive of perceptual experience.[13] SME can hold that even if representations underpin conscious, personal-level phenomena—such as perceptual presence—this does not mean that the representations themselves should be construed in personal-level terms.

Extrapolating from this, one might further insist that the fear of 'infection' from representations at the subpersonal to the personal level character of experience is unfounded, or at least dubious. As Locatelli & Wilson (2017, p.199) note in their discussion of the debates between relational and representational accounts: 'Subpersonal contents are not contents of any conscious experiential state or episode,

[13] Alternatively, even if PP's generative model is taken to be fully constitutive of perceptual experience, some have argued that this model is not realised only in the brain, thus allowing for extended consciousness. Kirchhoff and Kiverstein provide detailed arguments for such a claim, telling us that: 'the generative model can be conceived as realised not only in the brain but also in an extended dynamic singularity: in a continuous dynamic network with internal and external feedback loops' (2019, p. 51).

and so do not qualify as contents of experience'.[14] If this applies to the representations implicated in perception, their role can remain fully subpersonal; perceptual experiences as a personal-level phenomenon need not themselves *represent* the world as being some or other way, despite the subpersonal involvement of representations.

In summary, if S-representations play a role in realising an agent's mastery over SMCs, this does not establish that perceptual experience itself should be characterised as 'representational'.

5.1 The reconstitution challenge

We have seen that subpersonal representations need not present a problem for sensorimotor relationalism if one either endorses deflationism (§4.2.2) or accepts SME's compatibility with a limited role for robust subpersonal representations (§4.2.3). However, there is another way that the apparently innocuous compatibility of mechanistic approaches and SME (highlighted in §3.2) might be threatened. This innocuousness relies upon the mechanist accepting the pre-established phenomenon set out by SME—i.e., a skill-laden interactive relationship with the world—as the appropriate object of investigation. Mechanistic explanation is, as discussed, well-placed to set out the implementational details of such a phenomenon, but it may also play a more expansive role. It may additionally specify, challenge, or reconceive the phenomenon. With regards to perceptual consciousness, mechanists may reject the possibility of defining a sufficiently clear personal-level target phenomenon without the input of mechanistic enquiry. Moreover, they may object to the fact that if mechanistic inquiry only sets out implementational details of SME's pre-established phenomenon, it plays little role with regards to supporting or falsifying the sensorimotor account. Where tension may arise then is if, in one's mechanistic explanation, one compromises the characterisation of perceptual experience offered by SME.

In offering mechanistic explanations, our understanding of the phenomenon sometimes evolves. Thus, although explanations begin by investigating the mechanism causally responsible for a phenomenon (every explanation must start somewhere), our conception of the phenomenon may shift as our investigation proceeds (cf. Bechtel & Richardson, 1993/2010; Herschbach & Bechtel, 2014; Kronfeldner, 2015; Bollhagem & Bechtel, 2022). Bechtel & Richardson (1993) refer to this as 'reconstituting the phenomenon'. For instance, as Bechtel & Richardson (1993/2010) discuss, investigation of the 'Mendellian trait' (the target phenomenon) was originally depicted in terms of a macroscopic trait, such as hair or eye colour. However, such traits were discovered to be something that was not governed by single genes. In response, the phenomenon being explained was re-characterised in terms of something which *could* be explained by a single gene acting as its mechanism, namely, enzyme activity (cf. Herschbach & Bechtel, 2014).

This history of memory research offers another example, pertinent to cognitive science: over the decades, researchers re-articulated the salient properties of the phenomenon they were seeking to explain. As Bechtel (2008) observes, memory research began

[14] For further helpful discussion of possibilities for reconciliation between representationalism and relationalism, and potential hybrid views, see Locatelli & Wilson (2017).

with the assumption that (a) memory is a single phenomenon, and (b) memories largely correspond with reality. However, research suggests two modifications to our standard view: that (a) memory is more than one distinct activity, and (b) memories are less veridical replications of reality and more active constructs that are deeply intertwined with other mental operations.

In discussing this trend (which they refer to as 'recharacterising the phenomenon'), Craver & Darden (2013) identify three possible 'mischaracterisations' of a phenomenon:

> (1) claiming that a phenomenon exists when there is none, (2) lumping together two separable phenomena produced by different mechanisms, and (3) incorrectly splitting one phenomenon into many. (p. 60)

In the case of memory, mistakes of both type (1) and (2) were committed to the extent that there is no robust operation corresponding to the veridical recapturing of past events and there are, in fact, multiple distinct capacities that were conflated. As it pertains to SME's characterisation of perceptual experience, (1) and (2) also seem the most pertinent, as we are not concerned with sub-categories of different perceptual phenomena but a seemingly singular phenomenon that concretely specifies veridical perceptual experience: skilful agent-environment interaction. In any case, the general point is that, for mechanism, our very conception of the phenomenon may change throughout an investigation.

It would thus seem that even if a mechanistic approach to perception begins with a characterisation of perceptual experience in terms set out by SME, reconstituting or recharacterisng the target phenomenon—what one takes oneself to be explaining when they're explaining perception—is possible. One's best construal of the phenomenon, from a mechanistic point of view, *may*, in line with the sensorimotor account, include parts of the environment as well as the embodied agent. However, if mechanistic inquiry is in the business of challenging and clarifying what the target phenomenon is, this will be contingent upon the phenomenon as SME defines it surviving mechanistic scrutiny. It is worth re-emphasising at this stage that SME, first and foremost, establishes itself as a theory of veridical perceptual *experience*. As discussed below, our folk notion of 'perception' might refer to several distinguishable phenomena.

For precision, let's set out two subtly different conclusions one might draw in relation to 'reconstituting the phenomenon' of perception: that (a) skilful agent-world interaction is a real phenomenon (hence, trivially, one can offer a mechanistic explanation of its underpinnings as per §3.2), but such a phenomenon is not to be identified with perceptual experience; or (b) skilful agent-world interaction does not correspond to any sufficiently precise, empirically tractable phenomenon—we cannot identify concrete working parts that correspond to this notion.

We can demonstrate how (b) might proceed with an analogy. Historically, many have unpacked the phenomenon of perceptual experience in terms of qualia—instances of private, ineffable 'what-it-is-likeness'. However, over the course of scientific investigation, some have concluded that qualia do not refer to anything real (e.g.,

Dennett, 1988). Indeed, as mentioned above, SME is intended to offer an approach to perceptual experience that forgoes qualia (e.g., see O'Regan & Noë, 2001). Thus, the phenomenon is reconstituted, eliminating its specification in terms of private, ineffable properties. Likewise, the thought goes, if one unpacks the phenomenon of perception in terms of 'skilful agent-world exploration' we might decide, over time, that nothing within our best evidence for the causal structure underlying perceptual experience corresponds to this notion. In this way, from a mechanistic perspective, there is no linear path from articulating what perceptual experience is to discovering its constituents.

One way of imagining the possibility of reconstituting the phenomenon of perceptual experience is through mutual manipulability, introduced above. Very roughly, if those undertaking a mechanistic investigation have reason to suspect that parts of the agent and the world are not mutually manipulable in appropriate ways relative to observations they take to be evidence for perceptual experience, then they must conclude that the phenomenon of perceptual experience cannot be constituted (mechanistically) by an agent-world interaction; in other words, they have reason to suspect SME's characterisation of the phenomenon is off the mark. For example, one's confidence in SME could be eroded if there were mounting experimental evidence that conscious perception-like experiences could be induced or significantly altered merely by manipulating neural structures without any agent-world interaction. One might conclude on such a basis that the phenomenon (perceptual consciousness) is constituted only by internal vehicles, hence, SME in its relationalist guise is false. Before reaching this conclusion, however, caution is required.

5.2 Responding to the reconstitution challenge

In responding to worries about reconstitution, first, it's worth stressing that the degree to which the mechanistic constituents of perception concur with SME's conception of perception is contingent; again, nothing within mechanism precludes *extended* mechanisms. Sensorimotor theorists can be seen as placing an empirical bet: it is a strength of SME that it guides empirical investigation of perception in a manner that can be strengthened or weakened by evidence i.e., it is falsifiable, or at least empirically accountable (see §3.1). Moreover, as we have seen, at least some aspects of SME appear congruent with at least some versions of predictive processing which is regarded by many as forming the front of our best scientific theory of perception. Thus, proponents of SME may take solace in current trends.[15]

Second, in reconstituting the phenomenon of perceptual experience, we ought to consider the possibility that 'perception' is polysemous. This means that 'perception'

[15] Evaluating the degree to which our current best empirical theories of perception allow for extended mechanisms is beyond the remit of this paper. However, even if SME were to ultimately lose this empirical bet, the approach could nevertheless provide a 'guiding heuristic', reminding us to consider how a target phenomenon (conscious perception) is related to sensorimotor features of cognitive agents. Whatever, exactly, the outcome of a mechanistic inquiry into perceptual experience is, SME can function as a useful corrective to brain-centric biases, asking us to consider the role of the body and its dynamic interaction with the environment.

encompasses SME's primary target, *conscious* perceptual experience, in addition to other sufficiently distinct phenomena. Thus, in seeming to reconstitute the phenomenon, one might shift to a *different* phenomenon, i.e., to something that is not the target of SME. In such instances, one has changed the subject rather than demonstrated SME is false. For example, one's mechanistic explanation may pick out something broader than conscious perceptual experience, say, more elementary sensory capacities of plants, bacteria, or artificial systems. Alternatively, in detecting some apparent inadequacy in SME in reflecting the target phenomenon, one ought to consider whether conscious perception is not one thing, but multiple phenomena. This latter option is important to consider since, if this were the case, SME might capture *some* phenomenon falling under the umbrella of conscious perception, but be more limited in scope than canonical versions suggest. In such instances, one has not changed the subject, but neither have they shown SME is false. The takeaway lesson is that mechanists must be careful in identifying *what* exactly they are offering a *how* explanation of, and remain open to lumping and splitting errors (for recent discussion on how to identify lumping and splitting errors in mechanistic explanation, see Wajnerman-Paz & Rojas-Líbano, 2022)

Third, there may be different senses of constitution, and these should not be conflated. Mechanistic constitution specifically concerns the differentiable parts and operations that comprise a phenomenon 'from the inside'. There may be other compatible notions of constitution. Mechanistic constitution does not, for instance, exhaust the necessary elements of a phenomenon. For example, certain conditions are necessary for neurons to fire action potentials, such as maintaining a variety of homeostatic states within acceptable parameters. However, the kidneys (which measure oxygen content) are not part of the mechanism for action potentials. If the enactive sense of constitution is broader, perhaps encompassing certain causally necessary (but not mechanistically constitutive) elements for conscious perception, then it may be compatible with the mechanistic constituents remaining internal to the perceiving organism (for sample discussion on enactivism and constitution, see Hurley, 2010; Herschbach, 2012; Gallagher, 2018; Kirchhoff & Kiverstein, 2019, Chapter 6; Jurgens & Kirchhoff, 2019; Lee, 2023). The sensorimotor theorist might, for instance, have in mind something like the notion of 'coupling' whereby two systems are coupled just in case the state variables of one are parameters of the other, and vice versa (e.g., Thompson, 2007). However, a phenomenon may be realised by a mechanism A that is coupled with another system B without B thereby being a mechanism component of the phenomenon. In short, mechanistic constitution is not necessarily the only notion of constitution, and so we should remain sensitive to the possibility that multiple, compatible notions are at play before inferring conflict between frameworks.[16]

Conclusion

[16] It's also worth acknowledging that mechanists are increasingly appreciative of the complexities surrounding individuating phenomena and their constitutive mechanisms, given how context sensitive cognitive capacities can be: an entanglement of important but non-constitutive elements such as background, inhibitory and modulatory conditions can affect the manifestation of a cognitive phenomenon, and prising apart the constitutive mechanism may be difficult in practice (for some discussion, see Wajnerman-Paz & Rojas-Líbano, 2022).

Mechanism offers an account of how cognitive science explains whilst SME offers an account of perceptual consciousness intended by proponents to be congruent with scientific explanation. This paper has explored the connections between SME and mechanism. Whilst SME and mechanism appear innocuously compatible, closer attention reveals tension in the form of the 'representation' and 'reconstitution' challenges. The representation challenge can be overcome to the extent that relationalism is abandoned, the representations in question are deflated, or we accept the limited role for robust representations in realising the skilful interaction between agents and worldly objects, which constitutes experience according to SME. We indicated the final option was the least problematic, if a robust form of SME is to be preserved. The reconstitution challenge is only pertinent to the extent our best mechanistic explanations do articulate the phenomenon of perceptual experience in a way that undermines the environmental constituents of experience implied by SME, which is not a given. This challenge may be welcomed by proponents of SME insofar as it clarifies the empirical conditions under which SME is supported.

Stepping back, it's clear SME can benefit those seeking mechanistic explanations of perception insofar as it intends a de-mystified and empirically tractable account of the nature of perceptual consciousness. In other words, it helps identify the phenomenon to be explained and guides the mechanist toward the right level at which a mechanistic explanation must top out (roughly, the scale of agent-environment interaction). At the same time, a mechanistic approach may benefit SME both by offering a means of empirically supporting or undermining the framework and in terms of setting out the implementational details of the phenomenon that SME has articulated.

Bibliography

Beaton, M. (2016). Sensorimotor direct realism: How we enact our world. *Constructivist Foundations*, 11(2), 265–276

Bechtel, W. (2008). *Mental mechanisms: Philosophical perspectives on cognitive neuroscience*. Taylor & Francis.

Bechtel, W., & Abrahamsen, A. (2005). Explanation: A mechanist alternative. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, *36*(2), 421–441.

Bechtel, W., & Abrahamsen, A. (2010). Dynamic mechanistic explanation: Computational modeling of circadian rhythms as an exemplar for cognitive science. *Studies in History and Philosophy of Science Part A*, *41*(3), 321-333.

Bechtel, W., & Richardson, R. C. (1993/2010). *Discovering complexity: Decomposition and localization as strategies in scientific research*. MIT Press.

Bielecka, K., & Miłkowski, M. (2020). Error detection and representational mechanisms. In J. Smortchkova, K. Dolega, & T. Schlicht (Eds.), *What are mental representations?* (pp. 287–318). Oxford: Oxford University Press.

Bollhagen, A., & Bechtel, W. (2022). Discovering autoinhibition as a design principle for the control of biological mechanisms. *Studies in History and Philosophy of Science*, *95*, 145–157.

Brewer, B. (2011). *Perception and Its Objects.* Oxford: Oxford University Press

Brooks, R. (1991). Intelligence without representation. *Artificial intelligence*, 47(1), 139–159.

Campbell, J. (2002). *Reference and Consciousness*. Oxford: Clarendon Press.

Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, *58*(1), 7–19.

Clark, A. (2015). Radical predictive processing. *The Southern Journal of Philosophy*, 53(S1), 3-27.

Chemero, A., & Silberstein, M. (2008). After the Philosophy of Mind: Replacing Scholasticism with Science. *Philosophy of Science*, *75*(1), 1–27. https://doi.org/10.1086/587820

Craver, C. (2007). *Explaining the brain: Mechanisms and the mosaic unity of neuroscience*. Oxford: Oxford University Press.

Craver, C. F., & Darden, L. (2013). *In search of mechanisms: Discoveries across the life sciences*. University of Chicago Press.

Degenaar, J. (2014). Through the inverting glass: first-person observations on spatial vision and imagery. *Phenomenology and the Cognitive Sciences*, 13(2), 373–393.

Dennett, D. C. (1988). Quining qualia. In A. Marcel & E. Bisiach, (Eds.) *Consciousness in modern science*. Oxford: Oxford University Press.

Dolega, K. (2017). Moderate Predictive Processing. In T. Metzinger & W. Wiese (Eds.), *Philosophy and Predictive Processing*. Frankfurt am Main: MIND Group. https://doi:10.15502/ 9783958573116.

Drayson, Z. (2021). Naturalism and the Metaphysics of Perception. In H. Logue & L. Richardson (Eds.), *Purpose and Procedure in Philosophy of Perception* (pp. 215–233). Oxford: Oxford University Press.

Fish, W. (2021). Perceptual Paradigms. In H. Logue & L. Richardson (Eds.), *Purpose and Procedure in Philosophy of Perception* (pp. 23–42). Oxford: Oxford University Press.

Fazekas, P., & Kertész, G. (2011). Causation at different levels: Tracking the commitments of mechanistic explanations. *Biology and Philosophy, 26*, 365–383.

Gallagher, S. (2018). New Mechanisms and the Enactivist Concept of Constitution. In M. P. Guta (ed.), Consciousness and the Ontology of Properties, London: Routledge

Gładziejewski, P. (2015). Explaining cognitive phenomena with internal representations: A mechanistic perspective. *Studies in Logic, Grammar and Rhetoric*, 40(53), 63–90.

Gładziejewski, P. (2016). Predictive coding and representationalism. *Synthese*, *193*(2), 559–582.

Gładziejewski, P., & Miłkowski, M. (2017). Structural representations: Causally relevant and different from detectors. *Biology & Philosophy*, 32(3), 337–355.

Glennan, S. (2002). Rethinking mechanistic explanation. *Philosophy of Science*, 69(3), 342–353.

Glennan, S. (2009). Productivity, relevance and natural selection. *Biology & Philosophy*, *24*(3), 325–339.

Glennan, S. (2017). *The new mechanical philosophy*. Oxford: Oxford University Press.

Herschbach, M. (2012). On the role of social interaction in social cognition: a mechanistic alternative to enactivism. *Phenomenology and the Cognitive Sciences*, *11*(4), 467–486. https://doi.org/10.1007/s11097-011-9209-z

Herschbach, M., & Bechtel, W. (2014). Mental mechanisms and psychological construction. In L. F. Barrett & J. Russell (Eds*.), The Psychological Construction of Emotion* (pp. 21–44). New York: Guilford Press.

Hurley, S. (1998). *Consciousness in action*. Cambridge, MA.: Harvard University Press.

Hurley, S. (2010). The varieties of externalism. In R. Menary (Ed.), *The extended mind* (pp. 101–154). Cambridge: The MIT Press. Bradford Books. https://doi.org/10.7551/mitpress/9780262014038.001.0001.

Hurley, S. L. and Noë, A. (2003). Neural plasticity and consciousness. *Biology and Philosophy* 18: 131–168.

Hutto, D., & Myin, E. (2013). *Radicalizing enactivism: Basic minds without content.* Cambridge, MA: MIT Press.

Jurgens, A., & Kirchhoff, M. D. (2019). Enactive social cognition: Diachronic constitution & coupled anticipation. *Consciousness and Cognition*, *70*, 1-10.

Kaplan, D. M. (2012). How to demarcate the boundaries of cognition. *Biology & Philosophy*, *27*(4), 545-570.

Kaplan, D. (2018). Mechanisms and dynamical systems. In S. Glennan & P. Illari (Eds.), *The Routledge handbook of mechanisms and mechanical philosophy* (pp. 267–280). New York: Routledge.

Kirchhoff, M., & Kiverstein, J. (2019). *Extended consciousness and predictive processing. A third-wave view*. London: Routledge.

Kirchhoff, M., & Robertson, I. (2018) Enactivism and predictive processing: a non-representational view. *Philosophical Explorations*, 21(2), 264-281.

Krickel, B. (2018). Saving the mutual manipulability account of constitutive relevance. *Studies in History and Philosophy of Science Part A*, *68*, 58-67.

Kronfeldner, M. (2015). Reconstituting phenomena. In U. Mäki, I. Votsis, S. Ruphy, G. Schurz (Eds.), *Recent developments in the philosophy of science: EPSA13 Helsinki* (pp. 169-181). Springer, Cham.

Lee, J. (2019). Structural representation and the two problems of content. *Mind & Language*, *34*(5), 606-626.

Lee, J. (2021). Rise of the swamp creatures: Reflections on a mechanistic approach to content. *Philosophical Psychology*, 1-24.

Lee, J. (2023). Enactivism Meets Mechanism: Tensions & Congruities in Cognitive Science. *Minds and Machines*, *33*(1), 153-184.

Leuridan, B. (2012). Three problems for the mutual manipulability account of constitutive relevance in mechanisms. *British Journal for the Philosophy of Science, 63*, 399–427.

Locatelli, R., & Wilson, K. (2017). Introduction: Perception Without Representation. *Topoi*, 36(2), 197-212.

Machamer, P., Darden, L., & Craver, C. (2000). Thinking about Mechanisms. *Philosophy of Science, 67*(1), 1-25.

Marr, D. (1982). *Vision*. New York, NY: W. H. Freeman.

Martin, M. (2004) The limits of self-awareness. *Philosophical Studies*, 121(1), 37–89.

Nanay, B. (2014). Empirical problems with anti-representationalism. In B. Brogaard (Ed.), *Does perception have content?* (pp. 39–50). Oxford: Oxford University Press.

Noë, A. (2004). *Action in perception*. Cambridge, Mass: MIT Press.

Noë, A. (2008). 'Reply to Campbell, Martin, and Kelly' *Philosophy and Phenomenological Research*, LXXVI (3), 666–673.

Noë, A. (2009). 'Conscious Reference', *Philosophical Quarterly*, 59(236), 470–482

Noordhof, P. (2021). Wading in the shallows. In H. Logue & L. Richardson (Eds.), *Purpose and Procedure in Philosophy of Perception* (pp. 191–214). Oxford: Oxford University Press.

O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely-moving rat. *Brain research*.

O'Keefe, J. (1976). Place units in the hippocampus of the freely moving rat. *Experimental neurology*, *51*(1), 78-109.

O'Brien, G., & Opie, J. (2004). Notes toward a structuralist theory of mental representation. In *Representation in mind* (pp. 1-20). Elsevier.

O'Regan, J. K. (1992). Solving the "real" mysteries of visual perception: the world as an outside memory. *Canadian Journal of Psychology*, 46, 461–488.

O'Regan, J. K. (2014). The Explanatory Status of the Sensorimotor Approach to Phenomenal Consciousness, and Its Appeal to Cognition. In J. M. Bishop, & A. O. Martin (Eds.), *Contemporary sensorimotor theory* (pp. 23-35). Heidelberg: Springer International.

O'Regan, J. K., & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. Behavioral and Brain Sciences, 24(5), 939-973.

O'Regan, J. K., Rensink, R. A., & Clark, J. J. (1999). Change-blindness as a result of "mudsplashes." *Nature*, 398(6722), 34–34.

Philipona, D. L. & O'Regan, J. K. (2006). Color Naming, Unique Hues, and Hue Cancellation Predicted From Singularities in Reflection Properties, *Visual Neuroscience*, 23(3–4), 331–9.

Piccinini, G. (2020). *Neurocognitive Mechanisms: Explaining Biological Cognition*. Oxford University Press.

Piccinini, G. (2022). Situated Neural Representations: Solving the Problems of Content. *Frontiers in Neurorobotics*, 16. https://doi.org/10.3389/fnbot.2022.846979

Ramsey, W. (2017). Must cognition be representational?. *Synthese*, *194*(11), 4197-4214.

Rensink, R. A., O'Regan, K. J., & Clark, J. J. (2000). On failures to detect changes in scenes across brief interruptions. *Visual Cognition*, 7(1-3), 127–145.

Rutar, D., Wiese, W., & Kwisthout, J. (2022). From representations in predictive processing to degrees of representational features. *Minds and Machines*. https://doi.org/10.1007/s11023-022-09599-6

Seth, A. K. (2014). A predictive processing theory of sensorimotor contingencies: Explaining the puzzle of perceptual presence and its absence in synesthesia. *Cognitive Neuroscience*, 5(2), 97-118.

Seth, A. K. (2015) The cybernetic Bayesian brain: from interoceptive inference to sensorimotor contingencies. In J. M. Windt, & T. Metzinger (Eds.) *Open MIND* (pp. 9–24). Frankfurt, Germany: MIND Group

Seth, A. K., & Tsakiris, M. (2018). Being a beast machine: the somatic basis of selfhood. *Trends in cognitive sciences*, *22*(11), 969-981.

Shea, N. (2018). *Representation in cognitive science* (p. 304). Oxford University Press.

Simons, D., & Chabris, C. (1999). Gorillas in our midst: Sustained inattentional blindness for dynamic events. *Perception*, 28, 1059–1074.

Thompson, E. (2007). Mind in Life: Biology, Phenomenology, and the Sciences of Mind. Cambridge, MA: Harvard University Press.

Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological review*, *55*(4), 189.

Vázquez, M. J. C. (2020). A match made in heaven: predictive approaches to (an unorthodox) sensorimotor enactivism. *Phenomenology and the Cognitive Sciences*, *19*(4), 653–684.

Wajnerman-Paz, A., & Rojas-Líbano, D. (2022). On the role of contextual factors in cognitive neuroscience experiments: a mechanistic approach. *Synthese*, 200, 402. Online first. https://doi.org/10.1007/s11229-022-03870-0

Walsh, K. S., McGovern, D. P., Clark, A., & O'Connell, R. G. (2020). Evaluating the neurophysiological evidence for predictive processing as a model of perception. *Annals of the New York Academy of Sciences*, *1464*(1), 242–268. https://doi.org/10.1111/nyas.14321

Ward, D. (2012). Enjoying the spread: Conscious externalism reconsidered. *Mind*, 121(483), 731–751.

Ward, D. (2022). Sensorimotor Relationalism and Conscious Vision*, The Philosophical Quarterly*, https://doi.org/10.1093/pq/pqac016

Williams, D., & Colling, L. (2018). From symbols to icons: The return of resemblance in the cognitive neuroscience revolution. *Synthese*, *195*, 1941-1967.

Wilson, R. A. (1994). Wide Computationalism. *Mind*, 103(411), 351–72.

Zednik, C. (2011). The nature of dynamical explanation. *Philosophy of Science*, *78*(2), 238-263. http://doi.org.10.1086/659221