
Radical Sensorimotor Enactivism & Predictive Processing

Providing a Conceptual Framework for the Scientific Study of Conscious Perception

Adrian Downey

In this paper I outline and defend a novel approach to conscious perception, which I label “*radical sensorimotor enactivism*”. The aims of the paper are two-fold: (1) to respond to a common objection to theories like radical sensorimotor enactivism— that they are empirically vacuous— and explain why, because radical sensorimotor enactivism uses (a non-representational version of) predictive processing to operationalize its sub-personal aspects, this objection cannot be levelled at the theory; and, (2) to argue that radical sensorimotor enactivism provides a better empirical account of conscious perception than predictive processing taken as a stand-alone theory. I conclude that radical sensorimotor enactivism provides one with a strong over-arching conceptual framework for the scientific study of conscious perception which clarifies the relation between existing strands of empirical work and provides practical guidance for future research. As such, it is worthy of further development, study, and application in empirical settings.

Keywords

Consciousness | Non-representational theories of cognition | Perception | Predictive processing | Sensorimotor enactivism

Acknowledgments

This paper (or variants thereof) was presented at the Postgraduate WIP conference at the University of Sussex, an E-Intentionality seminar at the University of Sussex, at the Situating Cognition: Agency, Affect, and Extension conference held at the University of Warsaw, and the Philosophical Aspects of the Predictive Processing Framework conference held at the Frankfurt Institute for Advanced studies. I would like to thank the respective audiences of each of these talks for helpful discussions and constructive critical feedback. Thanks as well go to Sarah Sawyer and Chris Mole, both of whom provided extremely valuable feedback on the ideas contained within this paper. For helpful feedback on the manuscript itself, I would like to thank Joe Morrison, Joe Dewhurst, and two anonymous referees. Finally, and most importantly, I would like to thank Thomas Metzinger and Wanja Wiese for both: 1) providing instructive feedback on this paper, and 2) facilitating the creation of this volume and organising the conference upon which it is based. A lot of hard work goes on behind the scenes in these kind of ventures, but although it may go unseen, it does not go unappreciated – thanks very much for all the help. This research was financed with funding from the AHRC.

1 Introduction

In this paper I introduce a novel approach to conscious perception, which I label “radical sensorimotor enactivism” (RSE). RSE is advanced within the intellectual tradition of ecological and enactive approaches to mentality. It is commonly objected that theories advanced within this intellectual tradition are incapable of explaining the brain’s role in conscious perception (Chemero 2009, p. 93; Clark 2009; Seth 2014). Consequently, it is thought that such theories provide mere descriptions of conscious perception which are empirically unilluminating. I explain how a non-representational version of predictive processing (PP) can be subsumed within RSE and used to empirically explain the brain-based aspects of this framework. By doing so, I show that one cannot reject RSE for failing to account for the brain’s role in conscious perception. Furthermore, I then argue that, not only can RSE account for the brain’s role in conscious perception, it actually provides a better account of the brain’s role in

conscious perception than PP taken as a stand-alone theory. RSE provides a powerful over-arching conceptual framework for the scientific study of conscious perception which helps to clarify and taxonomise existing strands of empirical work whilst providing guidance for future research. As such, I conclude that it is worthy of further research, development, and application in empirical settings.

My paper is structured as follows— in section two, I outline RSE. I explain that RSE is predicated on the *sensorimotor enactive* theory of conscious perception, but that it improves upon this theory because it can provide a better account of the sensorimotor enactive concepts of “*sensorimotor knowledge*” and “*attention*”. In section three I explain why theories like RSE are often thought to ignore the brain’s role in conscious perception, before outlining a non-representational version of PP and explaining how it can be subsumed within RSE. I show that this non-representational PP account can operationalise and empirically explain the brain-based aspects of RSE, and so it cannot be objected that RSE ignores the brain’s role in conscious perception. Finally, in section four, I provide two reasons for preferring this conjunction of RSE and non-representational PP over a representational PP account. I argue that RSE provides a better explanation of conscious perception because: it provides a better account of the inter-relation between the sub-personal, personal, and conscious levels of explanation; and, it can account for and better categorise a wider range of empirical work in cognitive science. Thus, I conclude that not only can RSE account for the brain’s role in conscious perception, it actually provides a better account of its role in conscious perception than rival theories (such as PP taken stand-alone).

2 What Is Radical Sensorimotor Enactivism?

In this section I outline a novel approach to conscious perception, which I label “radical sensorimotor enactivism” (RSE). I begin by outlining the *sensorimotor enactive* theory of conscious perception, upon which RSE is predicated. I explain that extant versions of this theory are problematic (from an enactive perspective) because the key concepts of “*sensorimotor knowledge*” and “*attention*” both require representation, or they are left entirely unexplained. Then, I provide a non-representational explanation of both of these concepts. Consequently, I arrive at a non-representational (or *radical*) version of sensorimotor enactivism, and thus at *radical* sensorimotor enactivism.

2.1 Sensorimotor Enactivism

Sensorimotor enactivism (O’Regan and Noë 2001; Noë 2004; O’Regan 2011) is a direct realist theory of (conscious) perception. Consequently, it takes (conscious) perception to involve a direct relation between the perceiving organism and its environment. It explains this relation to be enabled by the organism’s possession of *sensorimotor knowledge*, which is knowledge of the law-like relation between sensation and movement. For example, there is a law-like relation between an organism’s movements and its visual stimulation— when an organism moves closer to an object the object looms in the visual field, when it gets further away the object appears smaller, and so on. On sensorimotor enactivism an organism is thought capable of perceiving only when it *understands* this relation between sensory stimulation and movement. Finally, sensorimotor enactivism takes attention to be necessary for *conscious* perception. On this theory an organism will be conscious of its perceptual relation to the environment only when it attends.

Sensorimotor enactivism traces its intellectual roots directly from the non-representational tradition of enactive and ecological approaches to mentality (Gibson 1979; Varela et al. 1991; Ryle 1949/2000; cf. Chemero 2009, fig. 2.4). However, in spite of its non-representational heritage and its proponents’ arguments against representational theories of conscious perception (O’Regan and Noë 2001), extant versions of the view have not done enough to distance themselves from representation (as will be explained in the following sections). In particular, the key concepts of “*sensorimotor knowledge*” and “*attention*” are both explained in a manner which makes indispensable use of the concept

of “representation”. Consequently, these key terms have either received a representational explanation, or they are left underspecified. I am going to provide a non-representational account of both of them, and therefore arrive at a thoroughly non-representational version of sensorimotor enactivism.

2.1.1 Non-Representational Sensorimotor Knowledge

It is often emphasised that sensorimotor knowledge should be understood as a kind of practical skill, which is predicated upon the organism’s possessing a certain kind of embodied know-how. However, in spite of its purported non-representational credentials, the concept itself appears to require representation. This point has been most forcefully argued for by Hutto ([Hutto 2005](#); cf. [Hutto and Myin 2013](#), pp. 23-32). Hutto argues that, in order to amount to any kind of substantive claim about the nature of perception which can play any sort of explanatory role, the concept of “sensorimotor knowledge” requires representation. Consequently, in order to provide a properly *radical* version of sensorimotor enactivism, I must provide a thoroughly non-representational and yet explanatorily substantive account of sensorimotor knowledge.

At the sub-personal level of explanation, I propose that sensorimotor knowledge be explained in a psychological behaviourist manner (cf. [Block 2001](#)). Psychological behaviourists took mentality to be constituted by a series of relations between sensory input and motor output, with a given sensory input related to a given motor output (or series of motor outputs, depending on the organism’s phylogenetic and ontogenetic history). If we explain sensorimotor knowledge in this manner, then it can be taken to concern a series of relations between certain sensory inputs to, and certain other motor outputs from, the brain. On this sub-personal account of sensorimotor knowledge, the brain acts merely as a causal mediator between certain neural inputs and certain motor outputs, and so one avoids postulating representation at the sub-personal level (cf. [Ramsey 2009](#), ch. 4).

At the personal level of explanation, I propose that sensorimotor knowledge be understood in terms of Ryle’s theory of knowledge-how ([Ryle 1949/2000](#), ch. 2). On this theory, knowledge-how is understood wholly in terms of behavioural dispositions— if an organism possesses knowledge of how to perform an activity, then this knowledge-how will be exhibited in the organism’s actual behaviour and their counter-factual behavioural tendencies. For example, if an organism knows how to throw a stone, then it will exhibit stone-throwing behaviour in scenarios where stones are available and the organism deems it pertinent to throw them. If we apply this Rylean account of knowledge-how to personal level sensorimotor knowledge, then we should take the ability to perceive to be exercised when the organism is disposed to engage in perceptual behaviour. It will, for example, exhibit mating behaviour when it *sees* another member of its species of the requisite sex, when it *smells* the presence of such a member, and so on and so forth. If the organism is disposed to behave in a manner consistent with its understanding the law-like relation between sensation and movement, that organism can be considered to possess personal level sensorimotor knowledge. Once more, because this personal level account of sensorimotor knowledge is entirely concerned with dispositions, one definitively avoids representation.

By providing an account of sensorimotor knowledge upon which it is wholly constituted by non-representational causal mediation and/or dispositions at both the sub-personal and personal levels of explanation, we therefore arrive at a thoroughly non-representational account of sensorimotor knowledge. In section three, I will explain how this non-representational account can be made explanatorily substantive.

2.1.2 Non-Representational Attention

Although sensorimotor enactivism takes attention to be necessary for conscious perception, there has been hardly any work focused on providing a substantive account of what, exactly, is meant by “attention”. The only proposal within this vicinity has been provided by O’Regan, who argues that sen-

sensorimotor enactivism should adopt a *higher-order thought* approach to consciousness (O'Regan 2011). Given that attention is necessary for consciousness, it could therefore follow that attention is itself to be explained in terms of higher-order thought theory. Higher-order thought theory understands consciousness to occur when an organism possesses a higher-order thought about a lower-order mental state. For example, the organism's visual states will become conscious when it has a higher-order thought about those states. This theory indispensably requires representation (the higher-order thought is *about* the lower-order state).¹ Consequently, acceptance of a higher-order thought approach to attention requires acceptance of representation. Therefore, extant sensorimotor enactive accounts either do not explain attention at all, or they explain it in terms of representation. I am going to propose a non-representational theory of attention. First, I will outline Mole's metaphysical distinction between "*process*" and "*adverb*". Then, I will explain how Mole applies this metaphysical distinction to the case of attention. Finally, I will outline Mole's adverbial theory of attention, before providing a non-representational version of the theory and applying it to sensorimotor enactivism.

2.1.2.1 Distinguishing "Process" and "Adverb"

Mole's adverbial theory of attention is predicated upon a metaphysical distinction between the concepts of "process" and "adverb". He summarises this distinction as follows:

A taxonomy is a taxonomy on the basis of process if the taxonomy classifies events on the basis of having or gaining of a property *by an object*. A taxonomy is a taxonomy on the basis of manner if the taxonomy classifies events on the basis of the having or gaining of a property *by an event*. (Mole 2011, p. 29, italics in original)

In order to determine the metaphysical category of a given x , Mole argues that we must consider the following two questions (Mole 2011, ch. 2):

1. What is x ?
2. What does it mean for x to occur?

Mole argues that x should be accorded the metaphysical status of process if, in determining its metaphysical status, it is most natural to answer question (1) first. If, however, it is more natural to answer question (2) first, then Mole concludes that x should be accorded the metaphysical status of adverb. Mole uses the examples of "combustion" and "hastily" (a process and adverb respectively) to make this point clearer.

In order to determine the metaphysical category of combustion it is natural to answer question one first. Combustion occurs when an object gains the property of burning because, when an object combusts, a chemical reaction between oxygen and fuel occurs which results in burning. We can therefore classify combustion as a process which occurs to an object and causes it to gain the property of burning. Having arrived at an answer to question one, the answer to question two becomes obvious. For combustion to occur we require that the chemical process of burning occurs. For this reason, Mole labels combustion a *process-first* phenomenon— in order to determine what combustion is we have to first understand what the *process* of combustion is.

Consider now the adverb "hastily". This adverb can be used to describe many different types of event— the publication of a newspaper can be performed hastily, a person's walk to the train station can be performed hastily, and the actions of two particular Hobbits in Middle Earth can be hastily performed (or at least, considered as such from the perspective of a disapproving Ent). Each of these events involve entirely different processes and it would be difficult, if not impossible, to find a pro-

¹ Put more precisely, all extant versions of higher-order thought theory indispensably require representation. Although higher-order thought theory could perhaps be explained without representation, no such version of this theory has as yet been proposed.

cess which all of these hastily executed events have in common. Consequently, one cannot determine whether x was hasty solely by focusing on the process of x . Each of these events are similar because they are carried out in a similar manner (hastily), and not because they involve the same processes. As such, if we are to determine whether a given x was performed hastily, we need to understand the manner in which the process of x was carried out. Thus, given that it is most natural to answer question two first when determining whether an event was hasty, haste should be accorded the metaphysical status of adverb.

In summary— Mole argues that there is a metaphysical difference between certain processes and events, and he claims that this difference is captured by the metaphysical categories of “process” and “adverb”. Process-first phenomena can be grouped into the same set because they all involve an object gaining a property by undergoing a particular process. Adverbial phenomena, on the other hand, can be grouped into the same set because they all involve an event gaining a property by being carried out in a particular manner. In order to determine the metaphysical category of a given x , one must determine whether it makes more sense to ask first of x what process it involves, or to instead ask first of x the manner in which its process is carried out.

2.1.2.2 Applying the Distinction to Attention

In order to apply Mole’s metaphysical taxonomy to attention, we must ask the following two questions:

1. What is attention?
2. What is it for something to be done attentively?

(Mole 2011, p. 24)

Philosophers and psychologists have traditionally answered question one first. Consequently, they tend to believe that attention is a process-first phenomenon. The consensus view that attention is a process has not, however, led to a consensus opinion on what “attention” is. Indeed, there are so many different candidates for explaining the process of attention that most psychologists do not believe there is a single set of processes which are necessary and sufficient for attention. Mole argues against process-first theories of attention (arguments which I will not rehearse here), and in their stead offers his own positive proposal about how attention should be classified. In determining the metaphysical category of attention, Mole argues that it is most natural to ask question (2) first. Thus, he concludes that attention should be understood as an adverb.

According to Mole, when an organism attends to x , it does so by ‘attentively x -ing’. His *Cognitive Unison* theory proposes that attention occurs when an organism uses its cognitive resources in unison to perform a task. According to Mole, if an organism is to count as performing a task, the following three conditions must be met (Mole 2011, pp. 52-5):

1. The task must include the organism.
2. The organism must know-how to perform the task.
3. The organism must be putting their know-how to use.

Mole therefore argues that attention occurs when an organism uses its cognitive resources in unison to attend to task x . Organisms will only count as ‘attentively x -ing’ if it is the organism itself which performs x , the organism knows-how to x , and the organism is currently engaged in x -ing. If these three conditions are met, then the organism can be ascribed the adverbial property of ‘attentively x -ing’.

2.1.2.3 ‘Radicalising’ Cognitive Unison

Mole’s version of Cognitive Unison theory requires representation at both the personal and sub-personal levels of explanation. Mole takes attention to be a type of personal level cognition, and he argues that personal level cognition is representational:

A cognitive process, in this sense, is a process that operates on representations that encode their contents *for the agent of the task*: In order for a process to count as cognitive, there must be an agent-level contentful state whose content is directly determined, at least in part, by the content of the representations on which that process operates. (Mole 2011, pp. 57-58, *italics in original*)

Furthermore, at the sub-personal level of empirical implementation, Mole argues that Cognitive Unison is best understood as implemented via *biased-competition*:

If the cognitive unison theory gives us the correct account of *what* attention is, then the biased-competition model may give us the correct account of *how* many instances of this attention-realizing unison *come about*, and of how they get maintained. (Mole 2011, p. 133, *italics in original*)

Competition models of cognition take different areas of the brain to be engaged in constant competition with one another, with their end goal thought to be that of controlling neural processing. Biased-competition theory is a theory of attention which explains neural competition in terms of brain-based representations vying for control of the brain (Desimone and Duncan 1995). Consequently, biased-competition is a representational theory of attention, and so Cognitive Unison appears to require representation at the sub-personal level. Thus, in order to provide a non-representational version of Cognitive Unison, I must explain: how personal level cognition can be explained without representation; and, explain how biased-competition could occur without requiring representation.

Spatial constraints require that I do not argue for any particular non-representational theory of personal level cognition here. Instead, I am going to assume that a non-representational theory of personal level cognition *could* be given (see, for example Anderson 2014; Barrett 2011; Chemero 2009; Hutto and Myin 2013; Ramsey 2009).² Given that there are plenty of non-representational theories of cognition present in the literature, and given that Cognitive Unison is itself a type of cognition, it therefore follows that Cognitive Unison *can* be accepted without requiring representation. Therefore, provided one accepts a non-representational theory of cognition, one can arrive at a non-representational version of Cognitive Unison theory.

Where the empirical implementation of Cognitive Unison is concerned, although biased-competition theories are representational theories, there are non-representational alternatives to be found in the literature. One such approach is developed in the work of Anderson (Anderson 2014, esp. chs. 5 & 6; Anderson 2015). In his 2014 *After Phrenology*, Anderson argues for an extremely anti-modular, *neural-reuse* theory of cognition, upon which different neural regions are constantly altering their connections with one another in response to various task-demands. Anderson claims that brains are constantly engaged in the formation and dissolution of what he labels “transiently-assembled-local-neural-subsystems” (or TALoNS, for short). Although TALoNS can be studied via the traditional methods of systems neuroscience (for example Sporns 2010), Anderson argues that we can come to a better understanding of TALoNS if we adopt the *affordance competition hypothesis*.

“*Affordance*” is a theoretical concept which states that organisms perceive their environment in terms of the activities it affords (Gibson 1979). Humans, for example, are thought to perceive chairs

² Note, I am not assuming the truth of such non-representational theories of cognition. Rather, I am only claiming that non-representational theories of cognition *can* be provided. Even the staunchest representational theorist has to agree on this point, on pain of making representational theories of cognition true a priori and so empirically vacuous (Ramsey 2015).

as affording sitting behaviour, and so humans will perceive them as ‘sit-upon-able’. The Affordance Competition Hypothesis proposes that organisms are constantly engaged in processing multiple environmental affordances at any given time— affordances compete with one another within the brain, with the winning affordance being the one that ends up controlling behaviour (Cisek 2015).

Anderson argues that the brain’s neural dynamics should be studied in terms of affordance competition. By conceptualising neural dynamics in this manner, he contends that we can come to an understanding of why TALoNS are constantly being formed and re-formed ‘on the fly’.³ TALoNS are formed ‘on the fly’ because different affordances are competing with one another within the brain, with the winning affordance being the one which determines the nature of the TALoNS being deployed. Anderson labels his theory “*biased-affordance-competition*”, and he is quite explicit that it should not be understood in terms of representation. Because Anderson provides a non-representational version of biased-competition theory, we can apply his framework to Cognitive Unison theory and therefore arrive at a non-representational empirical implementation of it.

Cognitive Unison theory understands attention to be a type of personal level cognition which is empirically implemented via biased-competition. Although Mole’s own version of the view requires representation, I have explained how the view can be accepted without requiring representation— we accept a non-representational account of personal level cognition and a non-representational account of biased-competition (biased-affordance-competition). At this point, we have arrived at a non-representational (adverbial) theory of attention.

2.2 ‘Radical’ Sensorimotor Enactivism

Having explained how the key sensorimotor enactive concepts of “sensorimotor knowledge” and “attention” could both be accounted for without requiring representation, we can now provide a thoroughly non-representational version of sensorimotor enactivism. At this point, we therefore arrive at radical sensorimotor enactivism (RSE). RSE is predicated upon the sensorimotor enactive account of conscious perception. As such, it understands perception to be constituted by a direct relation between perceiver and environment which is enabled by the possession and exercise of sensorimotor knowledge. The organism becomes conscious of their perceptual relation to the environment only when they attend to it. RSE improves upon sensorimotor enactivism because: it provides a non-representational account of sensorimotor knowledge; and, it provides a non-representational account of attention.⁴

In order to better appreciate RSE, it will help to consider an example. Imagine the following scenario— two people who take the exact same route whilst walking through a park, and who carry out the same perceptual processes. One of the walkers, who I will label the “*mindful-walker*”, is walking through the park whilst practising mindfulness meditation. This being the case, they are consciously aware of their various perceptual relations to the environment as they make their way through the park. They can see the bright red of the rose in front of them, feel the cool breeze on their neck, and hear the gentling rustling of leaves behind them. The other walker, although taking the exact same route through the park, is not consciously aware of their perceptual relation to the environment. This *worrying-walker* is currently in the midst of a financial crisis, and they are trying to work out how best to extricate themselves from their distressing predicament. Because both walkers take the same routes and engage in the same perceptual processing, each walker can be said to perceive. This explains why neither walker trips over branches, why both divert their gaze when they (accidentally) look straight at

³ I derive my use of this phrase from Clark’s discussion of TALoNS (Clark 2016).

⁴ Of course, my claim— that RSE improves upon sensorimotor enactivism— is reliant on the prior assumption that non-representational theories of conscious perception are to be preferred over representational ones. Although I do think that non-representational theories of conscious perception should be preferred over representational ones (see Anderson 2014; Chemero 2009; Downey 2016, ch. 3; Hutto and Myin 2013), I will not rely on this claim here. Rather, I will simply re-iterate an earlier noted point: sensorimotor enactivism has been advanced primarily as a non-representational theory and draws its intellectual roots from the non-representational tradition of enactive and ecological approaches to mentality. Consequently, a non-representational version of the theory is to be preferred.

the sun, and so on and so forth. However, although both walkers are perceiving, only one walker is conscious of their perceptual states— the mindful-walker. RSE explains the difference between each walker in terms of attention: the mindful-walker is engaged in ‘attentive-perception’ whereas the worrying-walker is not. Whilst the former is using their cognitive resources in unison to attend to the task of perception, the latter is using their cognitive resources in unison to attend to their financial problems. Thus the difference between the two walkers, and the reason why one is conscious of their perceptual relation to the environment whilst the other is not, is an adverbial one. It concerns the manner in which each walker engages in the task of perception, and so is not solely concerned with the perceptual processing itself. Although both walkers are engaged in the same sort of perceptual processing, only one of these walkers (the mindful-walker) carries out the perceptual processing ‘attentively’.

3 RSE, Brains, and Predictive Processing

It is often objected that enactive and ecological accounts of conscious perception, such as RSE, are empirically inadequate. These theories are often considered to provide mere descriptions of conscious perception, whilst failing to provide any substantive empirical insights or interesting hypotheses. In particular, it is often argued that these theories are incapable of accounting for the (undoubtedly key) role of the brain in conscious perception (Chemero 2009, p. 93; Clark 2009; Seth 2014). The empirical paradigm of predictive processing, on the other hand, is generally considered to provide a good account of the brain’s role in conscious perception (Clark 2013; Clark 2016; Hohwy 2013; Hohwy 2016; Seth 2014). In this section, I am going to explain how predictive processing can be subsumed within RSE. Therefore, I will conclude that RSE can account for the brain’s role in conscious perception, and so RSE cannot be rejected for failing to do so.

3.1 What is Predictive Processing?

Predictive Processing (PP) conceives of the brain as a prediction machine whose sole role is to minimise prediction-error (Clark 2013; Clark 2016; Hohwy 2013). According to this theory, the brain uses Bayes’ theorem (or an approximation thereof) to update *internal generative models*, which are used to derive a best-guess as to the external causes of its current sensory input. The brain is constantly updating these models in response to *error-signals* (which mark a divergence between predicted input and actual input) and this occurs via the process of *active inference*. In active inference, the brain can either: change its model to fit the incoming sensory stream, or, change the incoming sensory stream to fit its model (for example, by moving). Perception, on this account, is constituted by the brain’s expectations as to the external causes of its perceptual input.⁵ Importantly, the terms “prediction” and “expectation” are wholly sub-personal and non-conscious. We, qua personal level agents, are not aware of the Bayesian nature of our perceptual processes. Rather, these perceptual processes are carried out at, and applicable to, the sub-personal level of explanation.

Given that the world we live in is inherently uncertain and unpredictable, it is inevitable that there will be errors within the brain’s internal generative models. In order to make the processing of error-signals manageable, the brain must have a mechanism for determining which error-signals should be scrutinised more carefully, and which should be ignored. PP theorists suggest that the brain can do this via assignments of *precision-weighting*. These assignments concern the amount of ‘weight’ or ‘gain’ which will be accorded to a given error-signal. The more weight or gain an error-signal receives, the more it will be able to influence the internal generative models created by the brain. Analogously, the lower the weight or gain assigned to a given error-signal, the less influence this signal is likely to have on the brain’s internal generative models. Precision-weighting thus provides the means through which the brain can make its task (of minimising prediction-error) manageable.

⁵ Note, these expectations are considered to be essentially *action-oriented* (Wiese and Metzinger 2017).

3.2 Non-Representational Predictive Processing

PP has been advanced and developed as a theory which is situated firmly within the cognitivist paradigm in cognitive science (Von Helmholtz 1867; Gregory 1980). As such, it is unsurprising that PP is generally taken to indispensably require representation (Clark 2016; Gładziejewski 2016; Hohwy 2013, ch. 8). In order to make PP compatible with RSE, I must provide a non-representational account of its key posits. Orlandi has recently provided such an account, arguing that the key PP posits of “prediction-signal”, “error-signal”, “prior probability”, “hyper-prior”, and “likelihood” all fail Ramsey’s *job-description challenge*. As such, she concludes that these posits do not deserve a representational status. In this sub-section, I will summarise Orlandi’s argument.⁶

In his 2009 *Representation Reconsidered*, Ramsey proposes a strategy for determining whether or not a given mechanism deserves to be described in terms of representation. In order to count as representational, Ramsey argues that the mechanism in question must play a causal and functional role in a system which is recognisably representational. He claims that we should determine whether or not its role is recognisably representational by comparing the properties of the mechanism with the properties of an artefact which we would pre-theoretically regard as representational. If the mechanism fulfils the same functional role as a proto-typically representational artefact, then that mechanism passes the *job-description challenge* and so deserves to be described in terms of representation.

Orlandi applies this argumentative strategy to PP and concludes that its posits fail the *job-description challenge* (Orlandi 2015; cf. Orlandi 2014). She first argues that the concepts “prediction-signal” and “error-signal” (which are often applied to low- and mid- levels of the neural hierarchy by PP theorists) are concerned only with proximal conditions and so can be explained in terms of causal covariation. She then argues that causal covariation accounts of representation are not truly representational because they fail the *job-description challenge*— these accounts concern only correlation between neural events, and such correlation is more naturally described in terms of causal mediation than in terms of representation (cf. Hutto and Myin 2013, ch. 4; Ramsey 2009, ch. 4). Consequently, she concludes that prediction- and error- signals in the brain should be described as non-representational causal mediators. Finally, she argues that the concepts of “prior”, “hyper-prior”, and “likelihood” themselves fail the *job-description challenge* because they act only as biases which pre-dispose the brain to enter into certain neural arrangements.⁷ Consider, for example, the functional role played by a pump in a typical water-fountain. When water reaches the bottom bowl of a fountain, it pools around the opening to a pump. This pump sucks a small proportion of water in and pushes it back up to the top of the fountain, where the water once more begins to trickle down. Only a small proportion of water will be pumped up to the top of the fountain at any given time, because if too much were to flow from the top at once the fountain would malfunction (e.g. it may overflow). The pump therefore plays the functional role of biasing the flow of a system of water, such that there will always be a large pool of water at the bottom of the fountain and only a small amount at the top. We would not, of course, describe the biasing function of this pump in terms of representation. According to Orlandi, higher-level aspects of the PP neural hierarchy are best described as playing exactly the same sort of biasing role within the brain. Just as the pump in a water-fountain plays the role of pre-disposing the fountain to instantiate a certain water-cycle when presented with water, so too do higher-level posits of PP play the non-representational function of pre-disposing neural systems to enter into certain arrangements whenever presented with a given environmental stimulus. Consequently, Orlandi concludes that PP processing is entirely non-representational.⁸

⁶ My thanks to Zoe Drayson, for bringing Orlandi’s work to my attention.

⁷ Orlandi provides a number of other arguments against treating higher-level PP processing in terms of representation, but spatial constraints require that I do not rehearse these arguments here.

⁸ Orlandi does, however, conclude that the results of this processing (the ‘winning hypothesis’) should be described in terms of representation. I explain why I think this conclusion should be rejected, in the context of RSE, in footnote 11 of this paper.

Although this conclusion may appear merely cosmetic or superficial— we replace the word “representation” with “causal-mediation” or “non-representational bias” when describing PP processing— it in fact has profound empirical consequences. If low- and mid- level aspects of PP processing involve mere causal mediation, then empirical work on PP has much more in common with pre-cognitivist paradigms in the mind sciences (such as psychological behaviourism) than is often realized (*cf.* Ramsey 2009). Furthermore, if higher-level PP posits are to be understood as non-representational biases, then this should lead to an emphasis on empirical frameworks which conceive of neural processing in such terms. We should therefore be led to emphasise empirical approaches such as *visual science statistics* (championed by Orlandi) or Anderson’s biased-affordance-competition framework (Anderson 2014; Anderson and Finlay 2014) because these approaches can be used to study and explain how non-representational biases within the brain are created and maintained. Thus, accepting a non-representational account of PP processing has consequences not only for our understanding of current empirical work, but also for the kind of empirical work we subsequently engage in.

In short, although PP is generally assumed to indispensably require representation, a closer look at the actual causal role its posits play within the neural hierarchy should lead one to the conclusion that PP does not require representation. Prediction- and error- signals should be understood in terms of non-representational causal mediation, whilst priors, hyper-priors, and likelihoods should be understood as non-representational neuronal biases. This conclusion has empirical consequences because it requires a re-conceptualisation of extant empirical work and favours certain approaches to future research. Having arrived at a non-representational account of PP,⁹ I am now going to explain how this account could be subsumed within RSE.

3.3 RSE and Non-Representational PP

In section two of this paper, we saw that RSE explains perception to be predicated on an organism’s possession of sensorimotor knowledge and conscious perception in terms of its ability to ‘attentively perceive’. Having stated that PP does account for brain-based processing, and having argued for a non-representational version of PP, I am now going to explain how this account can be subsumed within RSE. I will explain how non-representational PP can be used to provide an empirical explanation of the sub-personal aspects of RSE. Consequently, I will show that RSE *can* account for the brain’s role in conscious perception.

PP describes perception as a sub-personal process concerning action-oriented expectations about perceptual stimulation. As such, it matches exactly the description of sub-personal sensorimotor knowledge. PP can therefore be used to provide an operationalisation of brain-based sensorimotor knowledge (*cf.* Seth 2014). The reader should recall that sub-personal sensorimotor knowledge was explained to concern a series of relations between certain sensory inputs and certain other neural outputs. And we have already seen that PP processing should be understood in terms of causal mediation and biased neural processing. PP can therefore be used to explain the specific relations between certain neural inputs and certain other neural outputs. By using PP to explain a certain facet of neural perceptual processing and providing a non-representational account of its posits, one could therefore arrive at a fully worked out empirical explanation of the brain-based, sub-personal aspects of sensorimotor knowledge (as that concept is understood on RSE).

In addition to explaining sensorimotor knowledge, PP can also be applied to, and used to improve, Anderson’s biased-affordance-competition framework (Clark 2016, ch. 5). Clark notes that PP fits extremely well with Anderson’s framework because it too offers up an action-oriented view of perception which allows for extreme neural plasticity. Clark argues that PP improves upon Anderson’s framework

⁹ It is worth noting that Orlandi’s argument is controversial. Although I find her general line of argument convincing, many proponents of PP are likely to object to it (especially those who advocate so-called *conservative predictive processing* (Gładziejewski 2016; Hohwy 2013; *cf.* Clark 2015). Spatial constraints dictate that I do not delve further into the details of this debate here. However, the interested reader can consult (Bruineberg & Rietveld 2014; Chemero 2009; Downey unpublished; Ramsey 2009) for arguments germane to the one endorsed here.

because the concept of “precision-weighting” can be used to explain how the brain manages to configure (and re-configure) TALoNS on extremely rapid time-scales. He does so by noting that the influence a given neural area has on any other given neural area can be determined via precision-weighting. If signals from a particular area are given high precision, then they are likely to be propagated to other areas of the brain. If, however, they are given a low weighting, it is likely they will have very little influence within the neural system. Importantly, the precision-weighting assigned to a given signal is itself constantly in flux and can be rapidly altered. As such, precision-weighting allows for the rapid creation, and dissolution, of various neuronal coalitions. It provides the means through which different neural areas can be allowed to influence, or be prevented from influencing, one another.¹⁰ Thus, not only is PP compatible with Anderson’s framework, it can in fact improve upon it.¹¹

At the sub-personal level of empirical implementation, RSE explains perception in terms of sensorimotor knowledge and attention in terms of biased-affordance-competition. In this sub-section, we have seen that non-representational PP can be used to operationalise sensorimotor knowledge and that it can be used to improve Anderson’s biased-affordance-competition framework. Non-representational PP can, therefore, be used to provide an empirical explanation of the sub-personal aspects of RSE. By subsuming non-representational PP within RSE, we arrive at an empirically satisfactory explanation of the brain’s role in conscious perception. Thus, it cannot be objected that RSE ignores the brain’s role in conscious perception. In short— PP is ideally suited to play the (extremely important) role of explaining the sub-personal, brain-based aspects of RSE.

4 RSE Improves Upon Cognitivist PP

Thus far, I have outlined RSE and explained how applying a non-representational version of PP to the framework allows for an empirically adequate explanation of the brain’s role in conscious perception. In this section I am going to go further, and argue that RSE provides a better account of the brain’s role in conscious perception than PP taken as a stand-alone theory. Consequently, I conclude that not only *can* a non-representational version of PP be used to explain the sub-personal aspects of conscious perception on RSE. Such a use of PP is in fact *preferable*, because one arrives at a better account of conscious perception than that which can be provided by PP alone.

4.1 Benefit One— Explaining Levels of Explanation

One of the chief benefits of accepting RSE and combining it with non-representational PP is that one arrives at a straight-forward account of the relation between the sub-personal, personal, and conscious levels of explanation. RSE explains personal level perception in terms of behavioural dispositions. The categorical basis of these behavioural dispositions is then explained to be brain-based— organisms are capable of perceiving only if they possess brain-based sensorimotor knowledge.¹² This brain-based sensorimotor knowledge is itself understood to be constituted by a series of relations between certain sensory inputs and certain other motor outputs, and can be operationalised via a non-representational version of PP. Therefore, on RSE, the brain is considered to ‘give rise’ to personal level perception by controlling and driving behaviour. Conscious perception is then, similarly, explained entirely in terms

¹⁰ On a non-representational account of PP, precision-weighting should itself be explained as a (particularly fast-acting) neurochemically mediated bias.

¹¹ By conjoining biased-affordance-competition and PP, we are able to reject Orlandi’s conclusion that the results of PP processing are representational. Orlandi concludes that the ‘winning hypothesis’ should be considered representational because it fulfils her three criteria for representation: it is concerned with distal events; it is de-coupleable; and, it is used for the planning of organismal action. A conjunction of biased-affordance-competition and PP will identify the ‘winning hypothesis’ with the ‘winning affordance’. Because affordances are directly perceived by organisms, they neither concern distal events nor are they de-coupleable (at least, not in any interesting sense which would require representation, see Anderson 2014; Chemero 2009). Although the ‘winning affordance’ is used for the planning of organismal action, meeting this criterion alone is not sufficient for the ‘winning affordance’ to be ascribed a representational status (unless one deflates the meaning of representation so much that the concept becomes empirically vacuous, see Ramsey 2009; Ramsey 2015). Thus, on RSE, the ‘winning hypothesis’ should not be described in terms of representation.

¹² This account of the relation between sub-personal and personal levels of description is consonant with that outlined in (McDowell 1994). My thanks to an anonymous referee for helping me to clarify this point.

of behaviour. On this theory, conscious perception occurs when the perceiving organism's behavioural interaction with the environment is carried out 'attentively', with 'attentive perception' empirically implemented in the brain via biased-affordance-competition. Thus, RSE attributes to the brain the role of controlling behaviour, and it explains both perception and conscious perception entirely in terms of behaviour.

There is absolutely nothing mysterious or naturalistically unacceptable about the idea that a brain could control behaviour. Therefore, there is absolutely nothing mysterious about the relation between the sub-personal and personal levels of explanation on RSE. Furthermore, on this framework there is nothing mysterious about the relation between conscious and unconscious personal level perception. Personal level perception itself occurs when an organism exercises sensorimotor knowledge. If the exercise of this knowledge is performed attentively, then the organism's perceptual relation to the environment will become conscious. As such, the difference between conscious and unconscious perceptual processing is explained to be adverbial in nature, and there is nothing metaphysically suspicious or naturalistically awry with the existence of adverbial behaviour. Consequently, RSE provides a clear distinction between the sub-personal, personal, and conscious aspects of perception. Moreover, RSE makes sense of their existence and inter-relation without requiring any leaps of imaginative faith or speculative metaphysical theorising.

Importantly, RSE can provide this metaphysically innocuous construal of the sub-personal, personal, and conscious levels of explanation whilst providing a phenomenologically *compelling* account of (conscious) perception (Noë 2004; Ward 2012).¹³ Consider, in this vein, the sensorimotor enactive explanation of the difference between certain perceptual modalities. Each modality is considered to possess its own specific set of sensorimotor 'laws', which concern the law-like relation between movement and stimulation specific to the sensory modality in question. Visual sensorimotor 'laws', for example, concern the fact that objects will loom as we get closer to them, appear smaller as we move further away, and disappear from view if we close our eyes. In the case of audition, however, sound gets louder as we move closer to its source, it gets quieter as we move away, and closing one's eyes will have little (or no) effect on hearing. We can, therefore, provide a distinction between different modalities of perception by explaining the relevant modality-specific sensorimotor 'laws'.

Although these sensorimotor 'laws' are described entirely in terms of an organism's perceptual behaviour, they are also phenomenologically intuitive:

When it is brought to our attention that certain sensorimotor contingencies are characteristic of vision, others of hearing, others of touch, there is an 'aha!' response. (Hurley and Noë 2003, p. 146)

RSE is simply a non-representational version of sensorimotor enactivism. As such, accepting RSE allows one to explain (conscious) perception to be constituted entirely by behavioural interactions with the environment, whilst providing a phenomenologically intuitive account of (conscious) perception.

This presents an improvement upon PP, if taken as a stand-alone theory of conscious perception, because that theory faces the (quite familiar) problem of having to provide a phenomenologically compelling explanation of how sub-personal, non-conscious, brain-based processing could 'give rise' to personal level, conscious perception. Although this is an active area of research within PP, barring a complete conceptual or scientific revolution, it is difficult to see how any particular brain-based

¹³ It is on this point that RSE departs from, and improves upon, other theories which take conscious perception to be constituted by behaviour (Dennett 1993; Ryle 1949/2000; Wittgenstein 1953/2009). Unlike these behaviourist theses, which (arguably) fall prey to the problem of verificationism, RSE can provide a phenomenologically compelling account of conscious perception whilst taking conscious perception to be constituted by behaviour. The main reason that RSE avoids the problem of verificationism is that it presents one with a *dynamic* view of conscious perception, upon which the environment itself constitutively plays an *active* role in proceedings. For an extended argument for this point, the interested reader can consult (Downey 2016, ch. 3; cf. Clark and Chalmers 1998; Hurley 1998, p. 420-22; Hurley 2001).

account could provide a phenomenologically compelling explanation of the inter-relation between sub-personal, personal, and conscious levels of explanation. As Hurley and Noë explain:

By contrast, if it is brought to our attention that activity in a certain brain area is correlated with vision, we do indeed still want to ask: “But why does brain activity there go with what it is like to see, rather than to hear or touch?” (Hurley & Noë 2003, p. 147)

The problem with brain-based accounts is that, regardless of the specifics of the account in question, the identification of a certain brain-based process with a certain experience is always going to appear arbitrary. Furthermore, there is always going to be an air of mystery surrounding how the phenomenology of conscious perception could be contained within, or identified with, the brain. Even the most enthusiastic proponents of PP admit that this is a problem for the view, and they are quite forth-right in conceding that there has, as yet, been no concrete proposals as to how PP could provide a novel solution to it (Hohwy 2013, p. 202; Clark 2016, ch. 7, §16).¹⁴

In summary, RSE explains (conscious) perception entirely in terms of behavior (both organismal and environmental), and it does so whilst accounting for the phenomenology of conscious perception. Moreover, it provides this phenomenologically plausible account whilst providing an explanation as to how the sub-personal, personal, and conscious levels of explanation inter-relate without invoking any naturalistically mysterious or metaphysically suspect properties or posits. PP, when taken as a stand-alone theory of conscious perception, cannot provide such an account. RSE provides a more elegant and phenomenologically plausible account of the relation between different levels of explanation in conscious perception than PP taken alone. Therefore, RSE should be preferred as a theory of conscious perception on this basis.

4.2 Benefit Two— Accounting for Empirical Data

The second major benefit accrued by accepting RSE concerns its empirical consequences. RSE is advanced within the intellectual tradition of enactive and ecological approaches to conscious perception. PP, however, is advanced from within the intellectual tradition of cognitivist approaches to conscious perception. Although each of these traditions boasts their own empirically successful research programmes, they are generally thought to directly conflict with one another. Consequently, it is usually thought that acceptance of one approach requires a wholesale rejection of the other (Chemero 2009; Hohwy 2016). RSE does not, however, require a wholesale rejection of cognitivist insights. Quite the contrary, in fact— not only is PP compatible with RSE, it actually constitutes a crucial component of the overall RSE framework (as we have seen). Thus, by accepting RSE and taking non-representational PP to provide an explanation and implementation of the sub-personal aspects of that framework, we arrive at a theory which can take advantage of the empirical work carried out on behalf of both enactive/ecological and cognitivist traditions.

This is beneficial because both of these paradigms have given rise to empirically productive research programmes, which have led to numerous novel and predicted empirical results. Scientific research programmes survive largely on the basis of their empirical productivity, with empirical productivity itself generally thought to require an inference to the best explanation— the theory is empirically productive because it accurately describes its domain of study.¹⁵ Although one can reject even empirically successful research programmes (for example, because one thinks they do not correctly describe a given domain of study), in order to do so one must explain how the science can be successful despite presenting us with an incorrect model of the world. Because RSE can accept empirical work carried

¹⁴ For an extended argument for this point, the interested reader can consult (Downey 2016, ch. 3).

¹⁵ Of course, there is a large literature on inference to the best explanation (and the related topics of realism, instrumentalism, and anti-realism) within the philosophy of science. I do not intend to presume a definitive answer to questions within this topic here. Rather, I am simply making the point that, by and large, the success of a scientific research programme usually gives us (defeasible) reasons to accept realism about its posits.

out in both traditions, it can simply side-step this problem. There is no need to reject, or otherwise eliminate, the vast swathes of empirical work carried out in either scientific tradition. Thus, on RSE, we have no need to reject an inference to the best explanation in either domain of study.

In fact, not only can RSE accept empirical work carried out on behalf of two traditionally opposed scientific frameworks, it actually can be used to help illuminate the distinction between the two research traditions and provide guidance for future empirical research. RSE equates the sub-personal level with the brain, and explains the brain's role to be that of controlling behaviour. Furthermore, it champions non-representational PP as the theory which should be used to study the brain. Thus, if we are interested in investigating the sub-personal aspects of RSE, we can do so by applying the conceptual and empirical tools of non-representational PP to the study of the brain (see, for example, Hohwy *et al.* 2008).¹⁶ If, however, we are more interested in investigating the personal level of explanation, then we can do so by using the methods of enactive and ecological approaches to study the interaction, and inter-relation between, the organism and its environment (see, for example, Chemero 2009). Finally, if we wish to study conscious perception, we can simply study the manner in which the perceiving organism is able to perceive 'attentively' (Anderson 2014; Mole 2011). In short, RSE helps to demarcate between the different levels of explanation within the science of conscious perception and therefore provides guidance as to which tools and techniques are appropriate for a given area of study.

The second benefit of accepting RSE, then, is that one arrives at an empirical integration of enactive/ecological and cognitivist approaches to conscious perception. This is beneficial because theorists can then take advantage of the excellent empirical work carried out within both traditions. Furthermore, this conjunction provides a clear and clean conceptual distinction between empirical work on the sub-personal, personal, and conscious levels of perception that provides guidance for the methods and frameworks which scientists should be using to study a given aspect of conscious perception. Thus, not only does RSE help to simplify and clarify the conceptual terrain of empirical work, it also provides empirical guidance.

5 Conclusion

In this paper I have outlined RSE, explained how a non-representational version of PP can be used to empirically explain its sub-personal aspects, and argued that the resulting account of conscious perception is to be preferred over PP taken as a stand-alone theory. I began by outlining the sensorimotor enactive theory of conscious perception, and explaining that this theory is problematic because two of its key posits ("sensorimotor knowledge" and "attention") either require representation or are left explanatorily vacuous. I argued for an account of sensorimotor knowledge in which it is taken to be constituted entirely by non-representational causal mediation and/or behavioural dispositions. Then, I outlined a non-representational and adverbial theory of attention and argued that it should be applied to sensorimotor enactivism. As such, I arrived at a thoroughly non-representational version of sensorimotor enactivism, and so at *radical* sensorimotor enactivism.

It is often objected that theories such as RSE are empirically vacuous. In particular, it is often argued that these theories are incapable of accounting for the brain's role in mentality. I outlined a non-representational version of PP and explained how it could be used to empirically explain the sub-personal, brain-based aspects of RSE. Therefore, I concluded that RSE cannot be objected to on the basis that it ignores the brain's role in conscious perception. Then, I explained why RSE should in fact be preferred as an account of conscious perception over rival cognitivist theories (such as representational PP). I argued that RSE provides a better account of the inter-relation between the sub-personal, personal,

¹⁶ In their account of binocular rivalry Hohwy *et al.* accept a representational version of PP. In order to subsume this account within RSE, we would therefore have to accept a non-representational account of PP posits. We have already seen that such an account can be given. In the case of rivalry, it is the PP framework itself (and not its representational posits) which plays a key explanatory role (Anderson and Chemero 2013). Therefore, accepting a non-representational account of PP's posits would not weaken the explanation of rivalry provided by Hohwy *et al.*

and conscious levels of explanation than cognitivist theories, and that it can account for and guide a larger amount of empirical research.

RSE is a novel theory of conscious perception which provides an over-arching conceptual framework for the scientific study of conscious perception. It promises to unite a number of (seemingly incompatible) strands of empirical cognitive science whilst demystifying the very existence of conscious perception. In addition to taxonomising different areas of extant research and clarifying their scope and inter-relation, RSE provides guidance for the direction of future empirical work. The phenomenon of conscious perception has only recently been submitted to sustained scientific scrutiny. Although there has of late been an explosion of empirical work on the topic, the empirical evidence accrued vastly out-weighs our ability to taxonomise and understand it. RSE provides a framework which can help to simplify this task substantially. I therefore conclude, on this basis, that it is worthy of further research, development, and critical scrutiny.

References

- Anderson, M. (2014). *After phrenology: Neural reuse and the interactive brain*. Cambridge MA: MIT Press.
- (2015). Précis of after phrenology: Neural reuse and the interactive brain. *Behavioural and Brain Sciences*, 1-22.
- Anderson, M. L. & Chemero, T. (2013). The problem with brain GUTs: conflation of different senses of “prediction” threatens metaphysical disaster. *Behavioral and Brain Sciences*, 36 (3), 204–205.
- Anderson, M. & Finlay, B. (2014). Allocating structure to function: The strong links between neuroplasticity and natural selection. *Frontiers In Human Neuroscience* (7), 918.
- Barrett, F. (2011). *Beyond the brain: How body and environment shape animal and human minds*. New Jersey: Princeton University Press.
- Block, N. (2001). Behaviourism revisited. *Behavioural and Brain Sciences*, 24 (5), 977- 978.
- Bruineberg, J. & Rietveld, E. (2014). Self-organization, free energy minimization, and optimal grip on a field of affordances. *Frontiers in Human Neuroscience*, <http://dx.doi.org/10.3389/fnhum.2014.00599>.
- Chemero, A. (2009). *Radical embodied cognitive science*. Cambridge, MA: MIT Press.
- Cisek, P. (2015). Cortical mechanisms of action selection: The affordance competition hypothesis. *Philosophical Transactions of the Royal Society of London B (Biological Sciences)*, 362 (1485), 1585–1599.
- Clark, A. (2009). Spreading the joy? Why the machinery of consciousness is (probably) still in the head. *Mind*, 118 (472), 963-993.
- (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioural and Brain Sciences*, 36 (3), 181- 204.
- (2015). Radical predictive processing. *The Southern Journal of Philosophy*, 53 (1), 3-27.
- (2016). *Surfing uncertainty: Prediction, action, and the embodied mind*. New York: Oxford University Press.
- Clark, A. & Chalmers, D. (1998). The extended mind. *Analysis*, 58 (1), 7-19.
- Dennett, D. (1993). *Consciousness explained*. London: Penguin Classics.
- Desimone, R. & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, 18, 193-222.
- Downey, A. (2016). *Radical sensorimotor enactivism*.
- (unpublished). Predictive processing and the representation wars: A victory for the eliminativist (via fictionalism). *Manuscript*.
- Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston: Houghton Mifflin.
- Gregory, R. (1980). Perceptions as hypotheses. *Philosophical Transactions of the Royal Society of London B (Biological Sciences)* (290), 181-197.
- Gładziejewski, P. (2016). Predictive coding and representationalism. *Synthese* (193), 559.
- Hohwy, J. (2013). *The predictive mind*. Oxford: Oxford University Press.
- (2016). The self-evidencing brain. *Noûs*, 50 (2), 259–285.
- Hohwy, J., Roepstorff, A. & Friston, K. (2008). Predictive coding explains binocular rivalry: An epistemological review. *Cognition*, 108 (3), 687–701.

- Hurley, S. (1998). *Consciousness in action*. Cambridge, MA: Harvard University Press.
- (2001). Perception and action: Alternative views. *Synthese*, 129, 3-40.
- Hurley, S. & Noë, A. (2003). Neural plasticity and consciousness. *Biology and Philosophy*, 18 (1), 131-168.
- Hutto, D. (2005). Knowing what? Radical versus conservative enactivism. *Phenomenology and the Cognitive Sciences*, 4, 389-405.
- Hutto, D. & Myin, E. (2013). *Radicalizing enactivism: Basic minds without content*. Cambridge, MA: MIT Press.
- McDowell, J. (1994). The content of perceptual experience. *The Philosophical Quarterly*, 44 (175), 190-205.
- Mole, C. (2011). *Attention is cognitive unison: An essay in philosophical psychology*. Oxford: Oxford University Press.
- Noë, A. (2004). *Action in perception*. Cambridge, MA: MIT Press.
- O'Regan, J. K. (2011). *Why red doesn't sound like a bell: Understanding the feel of consciousness*. Oxford: Oxford University Press.
- O'Regan, J. K. & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioural and Brain Sciences*, 5, 939-73.
- Orlandi, N. (2014). *The innocent eye: Why vision is not a cognitive process*. Oxford: Oxford University Press.
- (2015). Bayesian perception is ecological perception. *BrainsBlog*, 1-19.
- Ramsey, W. (2009). *Representation reconsidered*. Cambridge: Cambridge University Press.
- (2015). Must cognition be representational? *Synthese*, 1-18.
- Ryle, G. (1949/2000). *The concept of mind*. Chicago: Chicago University Press.
- Seth, A. (2014). A predictive processing theory of sensorimotor contingencies. *Cognitive Neuroscience*, 5 (2), 97-118.
- Sporns, O. (2010). *Networks of the brain*. Cambridge, MA: MIT Press.
- Varela, F., Thompson, E. & Rosch, E. (1991). *The embodied mind: Cognitive science and human experience*. Cambridge, MA: MIT Press.
- Von Helmholtz, H. (1867). *Handbuch der physiologischen Optik*. Leipzig: Leopold Voss.
- Ward, D. (2012). Enjoying the spread: Conscious externalism reconsidered. *Mind*, 121 (483), 731-751.
- Wiese, W. & Metzinger, T. (2017). Vanilla PP for philosophers: A primer on predictive processing. In T. Metzinger & W. Wiese (Eds.) *Philosophy and predictive processing*. Frankfurt am Main: MIND Group.
- Wittgenstein, L. (1953/2009). *Philosophical investigations*. London: Wiley-Blackwell.