

The Body as Laboratory

Prediction-Error Minimization, Embodiment and Representation

Christopher Burr (University of Bristol)

Max Jones (University of Bristol)

Citation: Christopher Burr & Max Jones (2016) The body as laboratory: Prediction-error minimization, embodiment, and representation, Philosophical Psychology, 29:4, 586-600, DOI: 10.1080/09515089.2015.1135238

Abstract

In his (2014) paper, Jakob Hohwy outlines a theory of the brain as an organ for prediction-error minimization, which he claims has the potential to profoundly alter our understanding of mind and cognition. One manner in which our understanding of the mind is altered, according to PEM, stems from the neurocentric conception of the mind that falls out of the framework, and portrays the mind as “inferentially-secluded” from its environment. This in turn leads Hohwy to reject certain theses of embodied cognition. Focusing on this aspect of Hohwy's argument, we first outline the key components of the PEM framework such as the ‘evidentiary boundary’, before looking at why this leads Hohwy to reject certain theses of embodied cognition. We will argue that although Hohwy may be correct to reject specific theses of embodied cognition, others are in fact implied by the PEM framework and may contribute to its development. We present the metaphor of the ‘body as a laboratory’ in order to highlight what we believe is a more significant role for the body than Hohwy suggests. In detailing these claims, we will expose some of the challenges that PEM raises for providing an account of representation.

Introduction

In a recent paper titled *The Self-Evidencing Brain*, Jakob Hohwy (2014) outlines his theory of the brain as an organ for prediction-error minimization (PEM), which he claims has the potential to profoundly alter our understanding of perception, action, attention and further aspects of cognition. The theory states that the brain seeks to maintain accurate models of the body and the environment, by predicting incoming sensory data in a top-down hierarchical manner (cf. Clark, 2013b; Friston, 2010; Hohwy, 2013, for summaries).

In this paper we restrict ourselves to exploring the commitments of PEM to the mind-world relation, and specifically Hohwy's claim that, “PEM should make us resist conceptions of this relation on which the mind is in some fundamental way open or porous to the world, or on which it is in some strong sense embodied” (Hohwy, 2014, p. 1).

We will argue that if we assume the validity of the PEM theory, then certain aspects of Hohwy's arguments *against* embodied cognition are correct. However, we will also argue that certain aspects of embodied cognition, such as its emphasis on the active and action-oriented nature of perception, are *implied* by PEM and may help further its development by bringing clarity to discussions concerning representation. In short, our argument can be seen as a direct response to the question that Hohwy poses at the start of his paper:

“How does a system such as the brain manage to use its *sensory input* to represent the *states of affairs in the world*?” (ibid., our emphasis)

The brain achieves this by utilizing *active sensorimotor predictions*, which have *high reliability*, in order to represent the world in an *action-oriented* manner. As such, the body should be understood as playing a far more significant role than Hohwy acknowledges.

We will begin by briefly outlining Hohwy's theory of prediction-error minimization, exploring the key concept of an *evidentiary boundary*, which commits him to a particular mind-world relation. We will then defend some of Hohwy's arguments against embodied cognition, whilst also highlighting potentially problematic consequences and omissions. Having addressed the arguments against embodied cognition, we will turn to the aspects of embodied cognition that are implied by the PEM approach. Taking these aspects into account has important implications for our understanding of both the role of the body and the nature of representation according to PEM. Inspired both by Gregory's (1980) interpretation of 'perceptions as hypotheses' and a more recent extension (Friston et al., 2012), we will defend a view where the body is best understood as the reliable and well-calibrated laboratory equipment that we use to probe the causal structure of the world. Given this view, which is compatible with Hohwy's framework, we finally turn to show why the debates concerning the nature of representation need to be revisited.

1 The Self-Evidencing Bayesian Brain

Whereas more traditional theories of cognition treat perception as a largely bottom-up process of incremental feature detection, PEM overturns this conception, instead placing an emphasis on top-down predictions about expected sensory data. These predictions emerge from multi-level, hierarchically-organized generative models, encoded as probability density functions, which are continuously modified by bottom-up error signals that in turn communicate mismatches between predictions and actual activity. This initial process is also accompanied by expectations of the precision of incoming sensory data (see Clark, 2015; Hohwy, 2013 for overviews; see Friston, 2010; Seth, 2013 for formal details).

The task faced by the brain, as presented by PEM, is to represent the states of affairs in the world, by maintaining accurate models that must generate their own evidence independent of any direct access to the world itself (see section 1.1). This evidence arises from the continuous flow of predictions and subsequent error-signals that serve to provide information to the brain as to which of its models are currently most accurate. Importantly, the hierarchical organization of the system is structured according to an increasing level of spatiotemporal scale, with the sensory receptors encoding input from the world at the lowest levels of the hierarchy.

The picture of the mind that falls out of this initial set-up is one that Hohwy (2014) describes as “neurocentric”. Anything outside of the brain is “inferentially secluded” from the internal models, and is treated as a “hidden cause” that must be inferred by the brain through a process that can be formalized using the tools of Bayesian statistics. What is shared by the vast majority of Bayesian approaches to cognition (e.g. Knill & Pouget, 2004; Tenenbaum et. al, 2011) is a commitment to modeling cognition as a form of hypothesis-testing. The motivation for this is clear when one considers the fact that the initial sensory input is consistent with a multitude of possible causes, and it's up to the brain to determine which is the most likely, given the current model. In other words,

perceptual input is vastly underdetermined with respect to the potential external causes, which it is supposed to represent. The proposed solution (following Helmholtz and Gregory), is to reduce the size of the hypothesis space by utilizing prior knowledge about the world, and a generative model of observing certain data given this prior knowledge, rather than relying solely on the sensory data itself. Such a task is well characterized using the formal methods of Bayesian statistics, and PEM builds on this framework, whilst retaining the core notion of the brain being a hypothesis-tester.

As a hierarchical framework, in PEM the posteriors at one level form the (empirical) priors for the level below, and are continuously shaped by the error signal that is propagated through the hierarchy. As Hohwy (2014, p. 4) states in the case of perception:

“Computationally, perception can then be described as empirical Bayesian inference, where priors are shaped through experience, development and evolution, and harnessed in the parameters of hierarchical statistical models of the causes of the sensory input.”

As briefly mentioned, the predictions are accompanied by an expectation of the precision of the sensory signal. This is important, as if a prediction has low uncertainty (high precision) over the range of possible states, it is more reliable, *ceteris paribus*, than a prediction that has high uncertainty (low precision). Whether the sensory signal is a reliable indication of the actual state of affairs in the environment, determines to what extent the models are updated. These sorts of precision expectations are further motivated by work from Bayesian approaches to cognition (especially in relation to attention, see Ernst & Banks, 2002; Feldman & Friston, 2010), and will play an important role in the later argument.

1.1 The Evidentiary Boundary and Free-Energy

According to Hohwy, the hidden causal structure of the world is always being inferred by the brain, from within the ‘Evidentiary Boundary’. As mentioned, the existence of this boundary entails a novel picture of the mind - one that is inferentially secluded from the hidden states of the world. In addition to the support this picture receives from the Bayesian Brain framework, the idea of an evidentiary boundary finds additional theoretical support, as well as a mathematical generalization, from a theory known as the free-energy principle.

The free-energy principle states that any (ergodic) self-organizing system, which can be described in terms of a Markov blanket, will appear to model and act on their world to preserve its functional and structural integrity. This unfolds in virtue of the minimization of an information-theoretic measure (free-energy) that bounds surprising sensory states for the system, and in turn leads to homeostasis (cf. Friston, 2010, 2013). It has been shown that the theory can provide a unifying account that bridges many disciplines (e.g. Bayesian inference, expected utility, information entropy and optimal control). It should also be noted that Hohwy (2014, 2015) has acknowledged the importance of the free-energy principle in providing theoretical support for the PEM account, as under simplifying assumptions, free-energy minimization can be understood as prediction-error minimization (Hohwy, 2013, p. 52). In what follows, many of the technical details have been omitted, and we refer the reader to key papers to further information.

In a Bayesian network, a Markov blanket is defined over a node X ; the set of nodes that comprise its parents; its children; and the other parents of all of its children. Any nodes in the network that fall outside the scope of the Markov blanket are independent of X , when conditioned on the set of nodes that comprise the Markov blanket. A Markov blanket thus creates a partition of states into

secluded inner states, and hidden external states, such that learning information about any of the external states will give no further information about the internal states. In the case of human agents the inner states consist of the states of the brain, separated from the external states of its environment by its perceptual states (parent nodes) and its active states (children nodes), which together constitute the Markov blanket (cf. Friston, 2013, for further details). The notion of a Markov blanket helps to make precise Hohwy's commitment to the "inferential seclusion" alluded to by the evidentiary boundary. The parameters explicit in the models of the brain are considered inner states, whereas the hidden states of the environment exist on the other side of the boundary that is induced by the Markov blanket. If Hohwy is correct in describing the brain, and thus the mind, in terms of a Markov blanket (or Evidentiary Boundary), then the mind must infer all external causes about the sensory signal that impinges upon it due to its secluded position.

Given this secluded position, the free-energy principle states that any (ergodic) self-organizing system must avoid *surprising* (considered in the information-theoretic sense) exchanges to ensure that their internal states remain within physiological bounds. As free-energy is minimized, the sensory input the agent receives (and subsequently explains away) becomes increasingly stronger evidence for the model in question (Friston, 2010).

Agents act in line with the predictions generated by their internal models in order to bring about less surprising perceptual states and remain within physiological bounds. By doing so, they implicitly maximize the Bayesian evidence of the internal generative models - if a model produces accurate predictions, it implicitly produces evidence for its own existence from within the Markov blanket (ibid.).

In section 3 we will discuss the important relationship between perception and action that falls out of this framework, but for now the important point to reiterate is the strict boundary that is implied by both the free-energy principle and PEM. Any agent that functions as these two accounts suggest, must infer the hidden states of the world from within the confines of an evidentiary boundary.

2 PEM and Embodied Cognition

Hohwy suggests that the PEM framework should make us resist conceptions of cognition as being in "some strong sense embodied" (Hohwy, 2014, p. 1) and that "the role of the body is real and substantial" but *only* when considered as being represented in the hierarchically-organized generative models of PEM (ibid., p. 17). However, it is important to be clear about the precise targets of Hohwy's arguments, since embodied cognition is not a single theory but a set of related yet distinct claims. It may be best understood as a research program, where the various strands are more closely united by their rejection of the prevailing cognitivist paradigm than by their mutual coherence (Shapiro, 2010, pp. 2-3). For present purposes, it is useful to highlight three distinct (though by no means exhaustive) theses of embodied cognition that are representative of Hohwy's own focus in (Hohwy, 2014):

1. Embodied Constitution of Mind: Cognitive processes are at least sometimes constituted by extra-neural parts of the body.

2. Radical Embodied Cognition: The best explanation of cognition does not invoke representational states. Instead it invokes organisms' active coupling with their bodies and environments.

3. Moderate Embodied Cognition: Cognition is best explained in terms of embodied representational states that utilize the same mental resources as are involved in perception and the guidance of action.

Hohwy explicitly rejects (1) and (2), whilst saying less about (3). In what follows we will argue that, given certain reasonable assumptions, Hohwy's rejection of (1) and (2) can be defended, but may provide Hohwy with less than he had hoped for. However, his neglect of (3) is somewhat surprising, given its apparent compatibility with, and significance for, the view that he is trying to promote.

2.1 Embodied Constitution of Mind

Hohwy is explicit in his rejection of views, such as (1), which suggest that the physical constituents of the mind extend beyond the brain and into the body. Proponents of this position have much in common with the extended mind hypothesis (Clark & Chalmers, 1998), however, cognitive extension is limited to parts of an organism's body, rather than also including parts of the external environment.

His denial of (1) rests on the assumption that PEM provides the correct description of mental *function* (i.e. maintaining accurate models for interacting with the world by way of the processes involved in prediction-error minimization). From this, he is able to reject the idea that the mind extends beyond the brain by utilizing the same kind of functionalist argument that his opponents use to argue for the contrary. PEM provides a more fine-grained definition of mental function, which allows Hohwy to be stricter than more liberal functionalists, who simply define mental function in terms of computation. For example, on a more liberal version of functionalism, one could argue that counting using one's fingers qualifies as a cognitive process that extends into the body, since the relevant finger positions serve as representations. However, since finger positions aren't hierarchically organized probabilistic representations, they fail to qualify as mental according to Hohwy's stricter definition of mental function.

Some form of cognitive extension is presumably possible, for example, if one were to hook up some form of silicon-based PEM device to the brain in the correct manner. However, this does nothing to detract from the fact that, according to PEM, every mental function that we know of is confined to the brain. Furthermore, this possibility of artificial cognitive extension does nothing to vindicate (1), which is concerned with whether natural parts of the body can fulfil the relevant function. It is important to note that this argument may only work against versions of (1) that are based on functionalist arguments. Whether similar arguments can be marshaled against proponents of embodied constitution who do not subscribe to functionalism (e.g. Menary, 2007) is an interesting open question that we are unable to address here. Given a functionalist approach and a definition of function in terms of PEM, we grant that Hohwy's neurocentric picture of the mind follows.

2.2 Radical Embodied Cognition

Hohwy rejects anti-representational versions of embodied cognition (2) (e.g. Chemero, 2011) on the basis that PEM is committed to a strictly representational approach to the mind. However, a lot hinges on how we are to interpret the representational commitments of the theory (see Hohwy,

2013, chapter 8, for a further exploration of the notion of representation in PEM, which focuses on statistical theories of representational content).

One way of interpreting the argument is to view the representational commitments as arising from the existence of *predictions*. As the idea of prediction seems to be an inherently representational phenomenon, PEM must therefore be committed to a representational approach. However, this can be challenged, as the term “prediction” can be used in two separate ways (Anderson & Chemero, 2013). Firstly, a prediction about a given event can take the form of a representation of its outcome. For example, the sentence “It will rain tomorrow” can be thought of as a prediction, and is also an explicit (linguistic) representation. Secondly, one can say that a certain event “predicts” another event when the two events are reliably correlated. For example, one might say that cows sitting in a field predicts the presence of rain and yet one might be hesitant to say that the cows' sitting is a representation. This latter sense of prediction need not be committed to any form of representation, and as such, Anderson and Chemero (ibid.) may be right to suggest that reference to “predictions” within the PEM framework is in this latter non-representational sense. Thus, if Hohwy's rejection of (2) were merely motivated by reference to predictions, then his argument would be flawed. However, there are other reasons for seeing the approach as committed to representations.

Another possible reason for seeing PEM as a representational theory is due to its commitment to hierarchical generative *models*. Modeling is uncontroversially a representational process, therefore, PEM is clearly committed to *some* form of representation. As such, it is tempting to take this as the motive for Hohwy's rejection of (2). However, it's important to note that many of the models PEM posits exist at specific levels in the hierarchy and represent the neural activity at the level below them (Friston, 2008). As a result, the overwhelming majority of this kind of representation is intra-neural. Only the most peripheral layers of the hierarchy directly model anything beyond the brain and these operate at extremely small spatial and temporal scales (Kiebel, Daunizeau, & Friston, 2008). As such, all that they can be said to represent are fleeting moment by moment impacts on small regions of our sensory receptors. It's difficult to see how they could represent the sorts of large-scale objects that typically populate our everyday discourse. As O'Regan and Degenaar (2014, p. 131) note:

“There may exist brain processes that can be viewed as hierarchically organized, with one layer functioning as if it “predicts” the activity of a lower layer. But this should not be taken to say that the higher levels represent external causes.”

Hohwy (2014, p. 15) acknowledges a thorny issue that arises from this conception of hierarchical modeling. In principle we could isolate the entire system minus the most peripheral layer, and we would still have a prediction-error minimizing system, complete with its own evidentiary boundary that separates it from the external world plus the peripheral layer. This process could be repeated, leading to a proliferation of nested hierarchical models (with their own evidentiary boundaries) representing both extra-neural and intra-neural processes. This is problematic for Hohwy's account, since the issue at hand is not whether there are any processes going on in the brain that could be described in representational terms, rather it is whether states of the mind can be understood as representations of states of the external world. Hohwy therefore privileges the entire brain as the object of study, as only when taken as a whole, can this system be plausibly described as representing the external environment. Hohwy can thus be seen to reject (2), as mental representation can be defined relative to this privileged evidentiary boundary.

This move leads to an uninformative notion of representation, in the sense that it doesn't provide a principled way of determining what parts of the external world the system represents or what parts or processes within the system serve as representational vehicles. All that we get is a definition of representation as some kind of relation between mental processes inside the evidentiary boundary and the world outside. In short, it provides us with a representation relation without providing any details about the *relata*.

It is important to highlight that this unfamiliar notion of representation, which emerges from the PEM account, is insufficient to help adjudicate in the debate between radical embodied cognition and the traditional computational accounts. As such, it will require careful handling in assessing whether Hohwy is right to reject (2) on the basis of representational commitments. We will return to this topic in section 5.

2.3 Moderate Embodied Cognition

Given Hohwy's rejection of anti-representational versions of embodied cognition, it is somewhat surprising that he makes no mention of more moderate strands of embodied cognition (3), which are explicitly committed to a representational approach to the mind (e.g. Barsalou, 1999; Clark, 1997; Prinz, 2002).

Proponents of moderate embodied cognition have argued that some form of mental representations need to be posited to account for so-called “representation-hungry” capacities for “off-line” cognition (Clark & Grush, 1999; Clark & Toribio, 1994). Furthermore, some proponents of embodied cognition accept the existence of representations, but insist that they are modal, in contrast with the amodal representations favored by classical cognitivists (e.g. Barsalou, 1999; Prinz, 2002). The debate about embodiment in this context is more a debate about what the neural vehicles of representation are, as opposed to a debate about whether they extend beyond the brain or whether such things exist at all. For example, Barsalou's simulation approach suggests that off-line cognition and, in particular, concepts, are constituted by reactivation of systems that are primarily dedicated to perception and the guidance of action (Barsalou, 1999, 2009).

There are a number of reasons why Hohwy's neglect of this major strand of embodied cognition is surprising. Firstly, proponents of (3) share his commitment to representations. Secondly, both proponents of (3) and of PEM share a commitment to dissolving the traditional boundaries between perception and action on the one hand, and cognition on the other. Thirdly, the major successes of PEM so far have been in explaining relatively low-level cognitive processes involved with interacting with the *immediate* environment. Certain versions of (3) could potentially be of benefit in extending the PEM framework to cover higher-level cognitive processes, such as off-line cognition, memory, long term planning, abstract cognition and imagination (Clark, 2013a). The notion of simulation, for example, has been argued to be closely related to PEM's central notion of prediction (cf. Barsalou, 2009).

Though Hohwy (2014, p. 17) may be correct when he claims that, “accommodating embodied cognition in this way happens within the strictures of the self-evidencing brain”, we believe that it is important to acknowledge the many gaps that PEM has yet to fill. A lot of important work is currently being done by proponents of so called moderate embodied cognitive science on the aforementioned representation-hungry capacities and other higher-level cognitive processes. Even if a PEM account can eventually explain these processes, we should not approach the work in a

myopic manner that may lead us to ignore vitally important contributions from outside the PEM framework - whether critical or supportive.

To develop this strand further, we now turn to explore active perception, a further aspect of the embodied approach, that is both compatible with and entailed by PEM. This is particularly relevant since understanding the notion of representation that emerges from PEM requires acknowledging the significant role of an agent's interaction with its environment. The ensuing discussion will thus demonstrate how the thesis of moderate embodied cognition can contribute to the further development of the PEM framework.

3 Interacting with the World

Having addressed the theses of embodied cognition that Hohwy rejects or neglects, it is important to address certain further aspects that are compatible with, and perhaps even inevitable, given the PEM perspective. Hohwy concedes that:

“The way perception unfolds will differ depending on the body's interactions with the world [...] In this sense, embodied cognition is inevitable, according to PEM.” (ibid., p. 16)

At first sight, this concession to embodied cognition can seem somewhat weak - even ardent critics of embodied cognition accept that the nature of perception depends on the body. However, once one pays closer attention to the PEM framework, and Hohwy's emphasis on “interactions”, it becomes clear that PEM gives a more substantial role to the body than initially appears. In particular, the PEM notion of *active inference*, suggests a dissolving of the distinction between perception and action. As such, the theory of active perception, a key aspect of embodied cognition, can be seen as an inevitable consequence of PEM and in turn leads to a far more significant role for the body than Hohwy acknowledges.

3.1 Active Inference

“Without the body, the system would only be able to minimize prediction error via passive perceptual inference and complexity reduction.” (ibid., p. 18)

It is possible within the PEM formalism to describe two distinct types of inference. *Perceptual inference*, involves updating models in light of incoming error signals, so that the driving sensory signals provide corrective feedback concerning the top-down predictions. *Active inference*, involves changing sensory input through action in order to match predictions. In active inference, the predictions are fixed, and used to drive action in order to fit the world to the model (cf. Adams, Shipp, & Friston, 2013). In effect, the two types of inference bring about the same result, and can be seen as re-descriptions of the same underlying imperative of prediction-error minimization (Friston, 2010).

Recent neuroanatomical evidence lends support to the blurring of perception and action (Shipp, Adams, & Friston, 2013). Traditionally, perception has been viewed as involving bottom-up transmission of information from sensory receptors to cognitive systems, whilst action primarily involves top-down transmission of motor commands to motor systems. As such, the traditional picture predicts anatomical asymmetry between perceptual and motor pathways. In actual fact, however, we find the opposite:

“The primary motor cortex is no more or less a motor cortical area than striate (visual) cortex. The only difference between the motor cortex and visual cortex is that one predicts retinotopic input while the other predicts proprioceptive input from the motor plant.” (Friston, Mattout, & Kilner, 2011, p. 138)

This lends support to the notion of active inference, which claims that action is also accounted for by a downwards cascade of predictive signals through motor cortex, to elicit motor activity, in much the same way as predictions descend through perceptual hierarchies.

PEM claims that both cascades predict sensory stimuli, whether it is proprioceptive, interoceptive or exteroceptive. Ascending pathways are also organized similarly in both perceptual and motor systems, with one exception being that certain prediction error signals are attenuated in the case of the latter. However, this is exactly what PEM predicts, since certain error signals need to be attenuated to allow for movement to take place, rather than updating of the model accompanied by immobility (Shipp, Adams, & Friston, 2013, p. 712).

Though initially helpful to separate perceptual and active inference as a heuristic for understanding the claims of PEM, it is misleading to equate *perceptual inference* with perception, and *active inference* with action. Perception is an active exploration of the agent's environment, and as such involves a continuous (and simultaneous) unfolding of *both* perceptual inference *and* active inference. Similarly, action involves both altering the environment by changing one's bodily state and monitoring these ongoing changes to update the model of one's own bodily state. As such, both perception and action, construed in folk psychological terms, involve a combination of *both* perceptual *and* active inference at the level of underlying cognitive processes.

The inclusion of active inference is taken to imply embodied cognition because the brain must do more than merely predict likely changes in the environment; it must predict the likely changes in sensory input that arise from the ongoing exploratory action of the agent. This interdependency of exploratory action and perception, which falls naturally out of the PEM framework, has been previously explored in the psychological and philosophical literature under the guise of active perception.

3.2 Active Perception

The theory of active perception, and the related sensorimotor theory, emerged from the tradition of ecological psychology (Gibson, 1979). The most important insight, in this regard, was to realize that there is more information available in the environment to an organism that is capable of active exploration than is available to a purely passive perceiver. Active perceivers are able to pick up on invariant features in the dynamics of sensory input as they explore their environment. For instance, as the eyes saccade from left-to-right the visual scene will shift from right-to-left in a predictable manner, relative to the speed and direction of saccadic-motion. An active perceiver can exploit regular relations between sensory input and motion of this kind in order to detect objective structural and causal features of the environment, which, on a traditional picture would need to be inferred from more basic sensory data. These predictable relationships between bodily movement and sensory input have come to be known as *sensorimotor contingencies* (O'Regan & Noë, 2001).

A further feature of the ecological approach to perception is to highlight the *action-oriented* nature of perceptual processes. On a more traditional theory of perception, information only becomes available for the guidance of action once a perceptual representation has been formed and passed

on to cognitive systems. However, this needn't be the case for active perceivers. This is because the kinds of complex invariant features that can be detected by an active perceiver have immediate relevance for action. Rather than first perceptually representing external objects and then inferring the consequences for action, active perceivers are able to perceive *affordances*, which are best understood as opportunities for action (Gibson, 1979).

The possibility of a direct link between perception and action has been used by proponents of radical embodied cognition to argue against representations (Chemero, 2011). However, it is possible to maintain the action-oriented nature of perception without taking the radical step of eliminating representations altogether. Instead, one can maintain that perception represents the world in an action-oriented manner (Clark, 1997; Mandik, 2005). As such, the seemingly representational nature of PEM is no reason to discount the potential significance of action-oriented perception. Furthermore, there is a certain sense in which, due to the tight coupling of sensory and motor processes, active and action-oriented perception come hand-in-hand. We perceive the world through active exploration of the environment and, in so doing, we perceive opportunities for further explorative activity.

The interdependency of perception and action in PEM due to the notion of active inference is something that many working within embodied cognitive science will be sympathetic to. However, one may still wonder about the “key role” that Hohwy (2014, p.16) assigns to the body. Though he considers the role of the body to be “real and substantial”, this is only when it's considered within the confines of the brain (i.e. as a hidden cause that must be inferred (and thus represented) by the hierarchical generative models) (ibid.). The following sections will now look at how this initial worry can be placated.

3.3 Extending PEM

In a recent paper, Seth (2014) provides a novel account of the way in which PEM can be extended to incorporate the aforementioned theory of sensorimotor contingencies (SMCs). Part of the motivation behind this is that possible neural or mechanistic implementations of sensorimotor contingencies remain unspecified, and that PEM can help operationalize the notion of SMCs.

As seen in the previous section, sensorimotor theory describes perception as a skillful activity, whereby an organism interacts with the world through mastery of the relevant laws that concern the contingent relationship between its body and the environment (O'Regan & Noë, 2001). Seth's (2014, p. 103) proposal is that hierarchically-organized generative models (HGMs) can be extended to account for this mastery. What is important to highlight about Seth's account is its emphasis on encoding *counterfactual* predictions within the generative models distributed throughout the cortical hierarchy. They are counterfactual in the sense that they don't merely make predictions about what probably will happen, but make predictions about various things that *would* happen conditional on an array of *possible* actions.

In doing this, Seth provides an operational account of the mastery of SMCs that is consistent with the PEM framework (e.g. how sensory inputs would change conditional upon a set of possible actions), and one which appears to fall naturally out of the considerations of active inference. However, there is an important distinction, which Seth draws attention to, that goes beyond the claims of active inference. Recall that in active inference, prediction errors are suppressed in order to resample the environment by acting on it, thereby aiming to bring about confirmatory signals for the respective prediction. In this sense, as Seth (ibid., p. 104) admits, the notion of counterfactual

predictions is already implicit in the dynamics of the priors predicting the sensory consequences of actions. However, Seth's account requires going further in encoding the counterfactual predictions *explicitly* as part of the priors in a HGM.

“That is, a counterfactually-rich HGM will model predicted future states (sensory signals, their external causes, and associated precisions) under a broad repertoire of different “controls” (those signals, not directly accessible to an agent, that cause movements).” (ibid.)

The explicit inclusion of these different “controls” (i.e. the subset of hidden states, separated by the evidentiary boundary, of which the agent has *indirect* control over) results in the models encoding potentially incompatible actions (e.g. whether to look right or left), in turn meaning that some predictions will inevitably be about situations that will never come to fruition. Determining which action will bring about the desired result, requires consideration of contextual knowledge, which will be determined by higher-levels.

As (Seth, 2015, p. 19) notes, there are a number of ways in which an experiment can provide evidence for a hypothesis; one can a) find evidence that will confirm a hypothesis, b) falsify or disconfirm it, or c) disambiguate between competing hypotheses by explaining away one or more of them. In the active inference framework, actions can be selected on the basis of any of these objectives. The use of counterfactual predictions provides a way for the system to compare hypotheses under this range of potential hypothesis-testing objectives. For example, if you wish to determine whether a sound is coming from your television or from outside, you could turn the volume up on the television to see whether the sound gets louder, or mute the television to test whether the sound persists. Depending on the actual state of affairs in the world, a combination of these actions (and possibly others) may be required. Being equipped with the active capacities to conduct this sort of hypothesis-testing is particularly useful when we have a way of determining which strategy is the most reliable and efficient.

3.4 Salient Experiments

In (Friston et al., 2012) the free-energy framework is adapted to explore the analogy of visual saccades as experiments. One of the examples given is a simulation where an image of a face is presented to an agent, and the responses (visual saccades) are considered ‘experiments’ that the agent performs to test a set of alternative hypotheses (i.e. an upright face, an inverted face, and a rotated face) (Ibid., p. 14). Testing the various hypotheses utilizes predictions concerning not only the expected sensory signal, conditional upon a certain action being performed, but also considerations about what they call the ‘saliency’ of a particular experiment. For instance, saccading to the right from a central area, based on the conditional belief that the ‘upright face’ model is the true hypothesis, will result in a mismatch if the ‘rotated face’ model is in fact the true state (under the former hypothesis the agent expects an eye, but will perceive a forehead if the latter hypothesis is true). The saccadic eye movements that the agent carries out can be understood using ‘saliency maps’ that are continuously updated as the belief in the model becomes more confident. If a saccade from a central area of the face is carried out, based on the expectation of sensing an eye, the saliency of this previously sampled location will be depleted. After all, performing the same experiment several times is unlikely to bring further evidence if the equipment is reliable and well-calibrated, and in ecologically valid situations waiting around to acquire further evidence can be costly. Furthermore, performing an action that results in the same evidence under two competing hypotheses, is less salient than a more discriminating experiment that could be performed in order

to explain away one of the hypotheses. These experiments show that an agent aiming to minimize prediction-error, must do so with considerations about how best to maximize the evidence for its own models. It is important to note that this is best achieved through the performance of salient experiments, and this requires that the agent encodes multiple counterfactual experiments, in order to compare their respective saliency. In this treatment it is important to highlight that saliency is not an attribute of the sensory cues, but rather an attribute of the (counterfactual) consequence of action.

In the case of real world scenarios, if the inherent uncertainty and noise of the environment is to be adequately dealt with, saliency cannot be determined without the additional expectations regarding the precision of specific predictions already detailed (see section 1). As we will see in the next section, our proposal for how a system such as the brain manages to deal with this problem, suggests a novel way of construing the problem of representation, which will bear on the discussions raised in section 2.

4 The Body as Laboratory

“Each movement we make by which we alter the appearance of objects should be thought of as an experiment designed to test whether we have understood correctly the invariant relations of the phenomena before us, that is, their existence in definite spatial relations.” (Helmholtz, 1971/1878 quoted in Friston, 2014)

In this section, we wish to take the analogy of ‘perceptions as hypotheses’ initially presented by (Gregory, 1980) and the subsequent development by (Friston et al., 2012) of ‘saccades as experiments’, one step further, by thinking of the body as a laboratory. Using this metaphor, the body of an agent can be seen as the reliable lab equipment that is used to carry out the aforementioned experiments, for the purpose of effective hypothesis testing. In scientific practice, hypotheses are tested using suitable lab equipment (e.g. a microscope is unlikely to be as useful for the study of acoustics as a microphone), and furthermore, effective hypothesis testing requires well-calibrated measuring devices in order to generate reliable data. By unpacking this metaphor, we hope to highlight a more vital role for the body than Hohwy acknowledges.

4.1 Sensorimotor Interactions and Reliability

“The perceptual and motor systems should not be regarded as separate but instead as a single active inference machine that tries to predict its sensory input in all domains: visual, auditory, somatosensory, interception and, in the case of the motor system, proprioceptive.” (Adams, Shipp, & Friston, 2013, p. 4)

We have already seen why the dynamics of our sensory and motor systems should not be considered as separate processes fulfilling different functions, rather than both fulfilling the ongoing pursuit of prediction-error minimization. This is not to reject the important distinction outlined earlier between perceptual inference and active inference. It is presumably possible to construct an artificial system that engages in purely passive perceptual inference. However, an important lesson of the theory of active perception is that, for organisms like ourselves, perception is never merely a process of passive perceptual inference - perception always involves an active exploration of the environment.

This is not because passive perception is impossible but because active perception allows us to access more information by utilizing the reliable and predictable bodily relations between motion and sensory input (cf. Gibson, 1979; O'Regan & Noë, 2001 for treatments related to this line of argument). By intervening on causal relations an agent can learn, and indeed shape the causal structure of its environment whilst testing its model. However, this does not have to be interpreted as a model of the world that the agent reconstructs separately from its interactions with it. Rather, controllability of one's actions is what allows an agent to intervene and test hypotheses in the first place, and should therefore factor into our understanding of what is being represented.

An apparent problem at this stage is that there seem to be robust patterns in the environment that do not immediately pertain to an agent's interactions, for example, the regular rising and setting of the sun. On this basis, it may seem like an agent should in fact encode within the generative models, representations of this interaction-independent causal structure. However, there are differences between our access to environmental and embodied regularities worth noting. Firstly, it is possible to decouple oneself from environmental regularities in a way that one cannot from bodily ones. Secondly, from the PEM perspective, it is the regularity of sensory patterns rather than environmental events that is significant due to the strict separation of mind and world by the evidentiary boundary. In the case of sunrises, the sensory input will vary depending on contextual features such as the direction one is facing, whereas sensorimotor contingencies are relatively invariant across contexts.

Sensorimotor interactions are more reliable because, unlike other statistical regularities in the environment, the agent can exploit them through action-oriented representations, which as some have argued could be adapted and reproduced over phylogenetic timescales (see Clark, 2013b; Friston, 2010, 2013 for some theoretical arguments in support of this claim). To put it another way, whilst the statistical regularities in the environment would have to be internalized through interactions and learning, it is likely that the statistical regularities pertaining to the ways in which our bodies interact with the environment have been stable enough over evolutionary time-scales so as to be genetically determined. It isn't necessary to learn about most important sensorimotor relationships because they can be built in to an organism's morphology and neural architecture (i.e. setting the priors in advance). Furthermore, the controllability of these interactions by the agent during ontogenetic development is likely to contribute significantly to the shaping of the representations.

We would expect that an agent is more likely to exploit the sorts of reliable organism-environment interactions that are contingent upon its phenotype, over less reliable, and more uncertain, organism-independent worldly structures. Interacting vicariously with the environment via the utilization of sensorimotor contingencies affords the agent a more reliable manner in which to minimize uncertainty. Just as scientists test hypotheses, by conducting experiments using well-calibrated lab equipment, so too perceivers must test their predictions by using their bodies to interact with their environment.

Hohwy argues that we are able to cope with noisy signals from the environment "because the world is a uniform kind of place that kindly affords reliable statistical inference" (2013, p. 224). However, as previously discussed, this reliability does not arise merely because the world is uniformly reliable. It arises precisely because certain parts of the environment, namely our bodies, behave in a more reliably predictable manner than the rest of the environment beyond them. The world would be a

far less kindly place if it weren't for the fact that our bodies are part of it and that their predictable behavior is, in some sense, under our own control.

We are now in a position to elaborate on our answer to the question posed at the start of this paper:

“How does a system such as the brain manage to use its *sensory input* to represent the *states of affairs in the world*?” (Hohwy, 2014, our emphasis)

The brain achieves this by utilizing *active sensorimotor predictions*, which have *high reliability*, in order to represent the world in an *action-oriented* manner. This is undoubtedly a different view of representation to more traditional conceptions, and so cannot, as Hohwy suggests, be seen to support a traditional representationalist account.

5 Representation in PEM

Under the view that has so far been detailed, the body is still represented by the mind, in accordance with Hohwy's argument and the notion of the evidentiary boundary. However, this does not entail that the body is represented in entirely the same way as features of the world.

Hohwy may be correct to emphasize that the body is outside the evidentiary boundary, but if the body was on the inside of the evidentiary boundary then it wouldn't play the special role of being a reliably predictable part of the environment. We are able to attain stability in a noisy and uncertain world because part of the environment (our body) is under our control and has been calibrated over evolutionary history to be able to generate reliably predictable data. Hohwy's claim, therefore, that “there is no difference between types of inference that rely on the body and types that don't” (ibid., p. 17) is misplaced. He may be right in the sense that both types of inference can be captured by the same underlying formal framework, however, there is a significant difference of degree in the reliability of the two processes. Although this difference is only a matter of degree, it is likely to be telling. Given that more reliable inferences that involve the body are available, it is unlikely that we utilize less reliable processes that don't. It is important to note that we are not claiming that inferences that don't involve the body are problematic or impossible. We are merely claiming that they are inferior, such that, as a matter of contingent fact, organisms are less likely to utilize such impoverished methods. We thus represent the causal structure of the world *vicariously*, by representing the possible results of exploratory actions.

Once one takes on board the consequences of the relationship between embodied cognition and PEM, it becomes clear that the notion of representation that falls out of PEM is very different from the standard notion of representation that anti-representationalists find troubling. The two most significant differences are that representations are, firstly, counterfactual and, secondly, organism-relative.

5.1 Counterfactual Representation

We have already seen why counterfactual predictions are required, but their relevance for the problem of representation is most clearly seen when one considers the problem that previous accounts of perceptual content attempted to solve. Traditional accounts describe the content of perceptual states as representing the world as it actually *is*, not as ways in which the world *will* likely be, conditional on actions that an organism *could* perform. Moving beyond a traditional account of

perceptual content, to a counterfactual, action-oriented one appears to be implied by PEM and the notion of active inference.

Seth's (2014) introduction of counterfactually-rich HGMs can seem like quite a radical departure from the original PEM framework, suggesting that we model possible and sometimes merely fictive eventualities. However, it is important to note that Hohwy's account is already a departure from a notion of objective representations of the way the world is. Perception, according to PEM, involves predictions about the way that the world is expected to be in the future. Therefore, it already involves a departure from the idea of perception as representing the world as it currently is.

The notion that mental representation is predominantly counterfactual is, without a doubt, extremely counter-intuitive. When introspecting on our own phenomenological experiences, it feels as though we are representing, or at least trying to represent, the single and objective way that the world actually is. It certainly doesn't feel as though we are representing the multiple ways that the world could be dependent on possible courses of action. However, it is important to acknowledge that the PEM framework is an attempt to explain the underlying mechanisms that give rise to phenomenal experience, rather than to explain phenomenal experience itself. Thus, if Seth's extension of PEM is correct, the experience of representing the way the world is, arises as a result of the underlying mechanism representing multiple ways that the world could be. As long as one is able to accept this distinction between phenomenal and representational content, the apparent counter-intuitiveness of counterfactual representations need not be a problem.

5.2 Organism-Relative Representation

An outcome of PEM's embrace of the active nature of perception, through its focus on active inference, is the inevitably action-oriented nature of representation. Organisms do not represent the world independently from their own interactions with it. Significantly, the ways in which an organism can interact with the world may vary considerably from species to species. Species with differing methods or abilities for interacting with the world could end up with significantly different perceptual content, even when placed in similar situations. Even further, this variability in perceptual content is likely to go beyond mere species-relativity, extending to organisms of the same species.

As is discussed in (Madary, 2015, p. 3), PEM suggests a notion of perceptual content such that "perceivers with different histories will have different predictions". This follows from the complex way in which the agent self-organizes, in response to the sensory signal from the environment, and in order to construct accurate models for future adaptive exchanges. As Clark (2015b, p. 2) acknowledges, the same sensory inputs could thus lead to very different perceptual states dependent on the way predictions have been altered in response to previous interactions.

If, in light of this, it still makes sense to understand perception in terms of representation then PEM "motivates an understanding of perceptual content that is always *organism-relative*" (Madary, 2015, 5, emphasis ours). Different organisms come equipped with different *lab equipment* to test the world in different ways, revealing particular aspects of the underlying casual structure that are relevant for their own specific capacities and goals.

The organism-relative nature of representation is particularly significant, since it challenges the traditional notion of an organism-independent categorisation of the world. On this account, which emerges from PEM, for example, we do not represent a chair as an objective organism-independent

entity. Rather, we represent it in an *action-oriented* manner, including the expected sensory consequences of our actions with respect to it. These consequences will be specific to us, and the type of organism that we are. We are in agreement with Clark (2015b, p. 5) when he states that in response to these conceptual shifts, it is difficult to see how we could capture the contents of such representations adequately using the terms and vocabulary of ordinary speech. However, this idea should not be taken as indicative of an idealist perspective, since the relations of an organism to its environment are still entirely objective features of the world, even if the ways in which it can interact with the world are organism-relative.

6 Conclusion

Returning to the comments made by Hohwy (2014) with regards to the explanatory scope of PEM, we can see that some of the conclusions concerning embodied cognition may not be as straightforward as initially assumed. Prior to PEM, the arguments that were made by anti-representationalists in favor of a more dynamic or enactive approach to cognition, had a clear target in the form of computational approaches to cognition. However this opposition needs to be readdressed in light of the PEM theory, which appears to require a novel and more nuanced description of what is meant by the notion of representation.

We believe Hohwy (ibid., p. 17) would be wrong to claim that PEM is able to accommodate embodied cognition within the strictures of the self-evidencing brain, whilst also supporting a traditional understanding of representation. Even if we grant that the body fails to play a constitutive role in cognition and that it is strictly speaking a part of the environment, a traditional picture fails to emerge. As we have argued, it is the body's role as a reliable and controllable part of the environment that enables the agent to engage in effective representation, and it is this special role that leads to a novel understanding of the nature of representation. Hohwy may be right in terms of eventually accommodating embodied cognition, but there is still much to do in terms of understanding the exact role that representations are playing in a PEM account of the mind, and embodied cognition may ultimately shape this discussion in ways that cannot yet be predicted.

The position defended in this article is that such an account must acknowledge the significant role that the body plays in shaping cognition, as well as the reliability of the sensorimotor interactions that embodied cognitive science investigates. The body-as-laboratory metaphor helps to bring these considerations into focus by highlighting the aspects of the PEM framework that imply the dissolving of the boundary between perception and action, as well as the organism-relative and action-oriented nature of representation. By demonstrating the compatibility between the PEM framework and work in embodied cognitive science, the body-as-laboratory metaphor may help further its development beyond the considerations of representationalism outlined here.

One such area where PEM requires development is in attempting to provide an account of off-line cognitive capacities, such as memory, reasoning, planning, imagination and abstract thought. As was mentioned earlier, Hohwy's neglect of moderate embodied cognition is somewhat surprising in this regard, since it seems both compatible with PEM and more suited to addressing these issues. Having acknowledged the organism-relative nature of representation on a PEM account, the need to turn to moderate embodied cognition becomes even more pressing. Both PEM and moderate embodied cognition suggest the dissolving of the boundaries between perception and action on the one hand

and cognition on the other. If this turns out to be empirically supported (see Anderson, 2014 for a promising treatment), then in addition to reconsidering the nature of representation, it may well turn out that PEM will impact areas beyond those detailed in the present article.

For the time being, PEM should not be seen as settling the debate about representation in favor of either traditional representationalism or radical embodied cognition - it may be too soon to call for peace in the war over the nature of representation (Clark, 2015b). Instead, Hohwy's work provides a fascinating new landscape in which the debate can proceed, invigorated by a rejuvenated notion of representation that differs significantly from that which was formerly held to be at stake.

Bibliography

Adams R., Shipp S. & Friston K. (2013). "Predictions not commands: active inference in the motor system", *Brain Structure and Function*, 218, 3, pp. 611–643.

Anderson, M. L. (2014). *After phrenology: Neural reuse and the interactive brain*. MIT Press.

Anderson M. L. & Chemero A. (2013). "The problem with brain GUTs: Conflation of different senses of prediction threatens metaphysical disaster", *Behavioral and Brain Sciences*, 36, 3, pp. 204–205.

Barsalou L. W. (1999). "Perceptions of perceptual symbols". *Behavioral and Brain Sciences*, 22, 4, pp. 637–660.

— (2009). "Simulation, situated conceptualization, and prediction", *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364, 1521, pp. 1281–1289.

Chemero, Anthony (2011). *Radical embodied cognitive science*, Cambridge MA, MIT press.

Clark A. (1997), *Being there: Putting brain, body, and world together again*, Cambridge MA, MIT press.

— (2013a). "The many faces of precision (Replies to commentaries on Whatever next? Neural prediction, situated agents, and the future of cognitive science)", *Frontiers in psychology*, 4.

— (2013b). "Whatever next? Predictive brains, situated agents, and the future of cognitive science", *Behavioral and Brain Sciences*, 36, 03, pp. 181–204.

— (2015a). "Embodied Prediction", *OpenMIND*, Ed. by T. Metzinger & J.M. Windt. Frankfurt am Main: MIND Group. url: http://open-mind.net/papers/embodied-prediction/at_download/paperPDF.

— (2015b). "Predicting Peace", *OpenMIND*, Ed. by T. Metzinger & J.M. Windt. Frankfurt am Main: MIND Group. url: http://open-mind.net/papers/predicting-peace-the-end-of-the-representation-wars/at_download/paperPDF.

— (Forthcoming). *Surfing Uncertainty: Prediction, Action and the Embodied Mind*, Oxford, Oxford University Press.

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

Clark A. & Grush R. (1999). "Towards a cognitive robotics", *Adaptive Behavior*, 7, 1, pp. 5–16.

Clark A. & Toribio J. (1994). "Doing without representing?", *Synthese*, 101, 3, pp. 401–431.

Ernst M. O. & Banks M. S. (2002). "Humans integrate visual and haptic information in a statistically optimal fashion", *Nature*, 415, 6870, pp. 429–433.

Feldman, H., & Friston, K. J. (2010). Attention, uncertainty, and free-energy. *Frontiers in human neuroscience*, 4.

Friston K. (2008). "Hierarchical models in the brain". *PLoS computational biology*, 4, 11

— (2010). "The free-energy principle: a unified brain theory?" *Nature Reviews Neuroscience*, 11,2, pp. 127–138.

— (2013). "Life as we know it", *Journal of The Royal Society Interface*, 10, 86, pp. 1–12.

— (2014). "Active inference and agency". *Cognitive Neuroscience*, 5, 2, pp. 119–121.

Friston K., Mattout J. & Kilner J. (2011). "Action understanding and active inference". *Biological Cybernetics*, 104, 1-2, pp. 137–160.

Friston K., Adams R. A., Perrinet L. & Breakspear M. (2012). "Perceptions as Hypotheses: Saccades as Experiments". *Frontiers in psychology*, 3.

Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*, Boston, Houghton Mifflin.

Gregory R. L. (1980). "Perceptions as hypotheses". *Philosophical Transactions of the Royal Society B, Biological Sciences*, 290, 1038, pp. 181–197.

Hinton G. E. (2007). "To recognize shapes, first learn to generate images". *Progress in Brain Research* 165, pp. 535–547.

Hohwy J. (2013). *The Predictive Mind*, Oxford, Oxford University Press.

— (2014). "The Self-Evidencing Brain". *Nous*, n/a–n/a.

— (2015). "The Neural Organ Explains the Mind". *Open MIND*. Ed. by T. Metzinger & J.M. Windt.

Frankfurt am Main: MIND Group. url: <http://open-mind.net/papers/the-neural-organ-explains-the-mind>.

Hohwy J., Roepstorff A. & Friston K. (2008). "Predictive coding explains binocular rivalry: an epistemological review", *Cognition*, 108, 3, pp. 687-701

Kiebel S. J., Daunizeau J. & Friston K. (2008). "A hierarchy of time-scales and the brain". *PLoS Computational Biology*, 4, 11.

Knill D.C. & Pouget A. (2004). "The Bayesian brain: the role of uncertainty in neural coding and computation". *TRENDS in Neurosciences*, 27, 12, pp. 712–719.

Madary M. (2015). "Extending the Explanandum for Predictive Processing". *OpenMIND*, Ed. by T. Metzinger & J.M. Windt. Frankfurt am Main: MIND Group. url: http://open-mind.net/papers/extending-the-explanandum-for-predictive-processing-a-commentary-on-andy-clark/at_download/paperPDF.

Mandik P. (2005). "Action-oriented representation". *Cognition and the brain: The philosophy and neuroscience movement*, Ed. A. Brook & K. Akins, New York, Cambridge University Press, pp. 284–305.

Menary R. (2007). *Cognitive integration: Mind and cognition unbounded*, New York, Palgrave Macmillan.

O'Regan K. J. & Degenaar J. (2014). "Predictive processing, perceptual presence, and sensorimotor theory". *Cognitive Neuroscience*, 5, 2, pp. 130–131.

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

Prinz J. (2002). *Furnishing the mind: Concepts and their perceptual basis*, Cambridge MA, MIT press.

Rao R. & Ballard D. H. (1999). "Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects". *Nature Neuroscience*, 2, 1, pp. 79–87.

Seth A. K. (2013). "Interoceptive inference, emotion, and the embodied self". *Trends in Cognitive Sciences*, 17, 11, pp. 565–573.

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

— (2015). "The Cybernetic Bayesian Brain". *OpenMIND*, Ed. by T. Metzinger & J.M. Windt. Frankfurt am Main: MIND Group. [url:http://open-mind.net/papers/the-cybernetic-bayesian-brain/at_download/paperPDF](http://open-mind.net/papers/the-cybernetic-bayesian-brain/at_download/paperPDF).

Shapiro L. (2010). *Embodied cognition*. New York, Routledge.

Shipp S., Adams R. A. & Friston K. (2013). "Reflections on agranular architecture: predictive coding in the motor cortex". *Trends in Neurosciences*, 36, 12, pp. 706–716.

Tenenbaum J. B., Kemp C., Griffiths T. L. & Goodman N. D. (2011). "How to grow a mind: Statistics, structure, and abstraction". *Science*, 331, 6022, pp. 1279–1285.