A Model Solution: On the Compatibility of Predictive Processing and Embodied Cognition

Luke Kersten[1] ⬤

Abstract

Predictive processing (PP) and embodied cognition (EC) have emerged as two influential approaches within cognitive science in recent years. Not only have PP and EC been heralded as "revolutions" and "paradigm shifts" but they have motivated a number of new and interesting areas of research. This has prompted some to wonder how compatible the two views might be. This paper looks to weigh in on the issue of PP-EC compatibility. After outlining two recent proposals, I argue that further clarity can be achieved on the issue by considering a model of scientific progress. Specifically, I suggest that Larry Laudan's "problem solving model" can provide important insights into a number of outstanding challenges that face existing accounts of PP-EC compatibility. I conclude by outlining additional implications of the problem solving model for PP and EC more generally.

1 Introduction

Predictive processing (PP) has emerged in recent years as a leading approach within the computational and cognitive neurosciences. In broad strokes, it maintains that a given system, such as a brain, creates and maintains a model of the causes of its sensory input. A system does not have direct access to the world but must infer hidden causes on the basis of sensory input and prior knowledge. PP is thought to deliver a simple yet compelling story for explaining a wide range of perceptual and cognitive processes and abilities, everything from vision and attention to consciousness and imagination (see, e.g., Clark 2013, 2016; Hohwy, 2013, 2016; Williams, 2019).

In a related vein, embodied cognition (EC) has experienced a similar steady rise to prominence within the cognitive sciences, albeit for different reasons. For this

✉ Luke Kersten
    luke.kersten@rub.de

1   Institute of Philosophy II, Ruhr University Bochum, Bochum, Germany

broad church, cognitive and perceptual processes are the result of the on-going and dynamic contributions of the body and world. Emphasising the action-oriented character of cognition, EC has proven a haven for those seeking an alternative to the reconstructionist and neuro-centric visions of cognition (see, e.g., Varela et al., 1991; Chemero 2009; Shapiro, 2011; Wilson & Foglia, 2017).

Recently, a number of authors have begun to wonder how compatible PP and EC might be. Clark (2016 p. 10), for instance, suggests that PP provides a home for the best of embodied and classical approaches, writing: "Predictive processing provides a meeting point for the best of many previous approaches, combining elements from work in connectionism and artificial neural networks, contemporary cognitive and computational neuroscience, Bayesian approaches to dealing with evidence and uncertainty, robotics, self-organization, and the study of the embodied, environmentally situated mind." For Clark, PP and EC are not only compatible, they form a unifying vision of cognition. Hutto (2018), on the other hand, worries that there is a tension looming between PP and EC. He writes, for instance: "PPC assumes that cognition is ultimately grounded in informational processing and the manipulation of representational contents. On this pivotal issue REC [radical enactive, embodied cognition] firmly disagrees with cognitivist versions of PPC [predictive processing cognition]. Put otherwise, when donning a cognitivist guise PPC is fundamentally at odds with REC" (p. 2448). For Hutto, PP's focus on internal models and probabilistic inference would appear to stand in tension with EC's emphasis on action and perception. Of course, these are only two choice examples, but they highlight a general issue that has occupied a number of authors.[1]

The task of sussing out PP-EC compatibility is an important one. Not only have PP and EC been heralded as "revolutions" and "paradigm shifts" in cognitive science, but they have also motivated a number of new and interesting areas of research, including work on vision (Noë, 2004, 2009), interoception (Seth, Suzuki and Critchley 2012), semantic representations (Meteyard et al., 2012), conceptual knowledge (Gallese & Lakoff, 2005), and religious experience (van Elk and Aleman 2017), to name only a few.[2] The outcome of the compatibility issue could have significant impacts on how we think about and study the mind. As Kirchhoff (2018a) has recently put the point: "making progress on this issue will no doubt yield substantial insights into the nature of mind" (p. 2342).

Given its importance, the current paper looks to weigh in on the issue of PP-EC compatibility. I argue that further clarity can be achieved by considering a model of scientific progress. Specifically, I suggest that Larry Laudan's "problem solving model" can provide important insights into a number of outstanding challenges that face existing accounts of PP-EC compatibility.

The paper unfolds in five parts. In Sect. 2 I begin by outlining two recent proposals on PP-EC compatibility: Clark (2015) and Hohwy (2018). Next, in Sect. 3, I outline three outstanding challenges facing PP-EC compatibility. These include: (i) how

---

[1] Other examples include Clark (2015, 2017), Hohwy (2016, 2019), Anderson (2017), Miller and Clark (2018), Kirchhoff (2018a, 2018b), and Venter (2021).

[2] For review, see Litwin and Miłkowski (2021) and Shapiro (2011).

to explain the theoretical status of PP and EC, (ii) how to identify the theoretical commitments of PP and EC, and (iii) how to make sense of the notion of compatibility at stake within discussion. Following this, in Sect. 4, I introduce the PSM, outlining three of its key components. This leads to the main proposal in Sect. 5. Here I argue that each of the three outstanding challenges can be successfully addressed using the PSM. Finally, in Sect. 6, I outline two additional implications of adopting the PSM for PP and EC more generally.

Two caveats are worth making before discussion gets going. First, I restrict analysis in what follows to the relationship between PP and EC.[3] Extended cognition, for instance, claims that an agent's environment or body can, on occasion, form part of its cognitive system. While EC sometimes makes constitution claims in this vein, extended cognition is generally seen as the more radical view in that it also includes parts of the environment (Wilson and Clark 2009; Kersten & Wilson, 2016; Kersten, 2017). Enactivism, moreover, conceives of cognition in terms of the biodynamics of living systems, often focusing on explaining intentionality and phenomenology via the active exploration of the environment, rather than the use of representations (Ward et al., 2017). While EC shares a similar focus on the action dynamics of the agent, it differs with respect to its target explanations and methodology. For instance, not only is EC generally thought to explain a wider range of phenomena than intentionality and phenomenology (Wilson & Golonka, 2013), but at least some formulations make space for talk of representations (Clark, 2008; Wilson, 2004). In focusing solely on embodiment, I am remaining neutral on a number of connected questions, such as the relation between PP and the question of life-mind continuity (Hutto & Myin, 2013, 2017) or how talk of prediction error minimisation fits with extended cognition (Kersten, 2022).[4]

Second, I do not directly address a number of prominent PP-adjacent ideas, such as the "free energy principle" or "active inference framework".[5] This is done for two reasons. The first is simply reasons of space; discussion of these elements would take analysis too far afield. The second is that a number of fruitful investigations have already been undertaken in these directions (see, e.g., Bruineberg et al., 2018; Constant et al., 2020).

2 PP-EC Compatibility

In this first section, I outline two prominent proposals on PP-EC compatibility. This discussion provides not only an important survey of existing work but it also sets the stage for the analysis to follow.

---

[3] For further discussion, see Kirchhoff and Robertson (2018), Kirchhoff and Kiverstein (2019), or Constant et al. (2020).

[4] Prediction error minimisation is the process whereby a system reduces the difference between what is predicted about the world and how the world actually is.

[5] The FEP is a mathematical description of principles of organisaiton for complex adaptive systems, while the active inference framework is a computational approach to understanding perception and cognition as a form of free energy minimisation.

2.1 Radical Predictive Processing

The first proposal comes from Clark (2015).[6] Clark's view, as mentioned, is that EC and PP are not only compatible but, in fact, form a unifying picture of cognition. To arrive at this position, Clark attempts to show that the key principles and insights of EC can be comfortably accommodated within PP.

Clark begins by surveying two key "lessons" from embodiment. The first is "productive laziness". Productive laziness describes the economical but effective use of strategies in problem-solving contexts when operating with limited time or processing power, such as relying on a trustworthy friend for restaurant recommendations. The second is cognitive scaffolding. Cognitive scaffolding describes an agent's use of the external world to offload computational task demands, such as how eye saccades minimise working memory demands during block-moving tasks (Ballard et al., 1997). Clark's suggestion is that, when taken together, these principles paint a dynamic picture of cognition, one in which embodied agents are constantly building, dissolving, and rebuilding temporary ensembles to exploit available resources, whether that be neural, bodily, or environmental.

Clark further points out that there are two interpretations available for PP. One is what he calls "Conservative Predictive Processing" (CPP). CPP adopts a "reconstructive" picture of prediction. According to this idea, inner models recapitulate the structure and richness of the real world for the purposes of planning, reasoning, and guiding action. The other is what he dubs "Radical Predictive Processing" (RPP). RPP adopts a "non-reconstructive" approach to prediction. This means that behavioural successes are not the product manipulating rich inner replicas but, rather, the result of action-perception cycles operating to keep sensory stimulations within certain bounds.

Clark suggests that PP is better understood along the lines of RPP rather than CPP. One of the reasons for this is that the latter reading conflicts with one of the core principles of PP: namely, that the "goodness of a predictive model is determined by accuracy minus complexity". In the context of PP models, this means that minimising complexity requires reducing computational costs as far as possible while performing a task.[7] For Clark, the trouble is that if PP models entail a reconstructive picture, then this would appear to contradict the "satisficing" principle. If PP systems are forced to learn and deploy the least complex solutions to accomplish a task, then they cannot be regularly constructing highly complex inner replicas of the world.

In contrast, the RPP reading offers a better fit with the lessons from embodiment. Clark returns to the previous examples to illustrate. First, because there is a simultaneous drive to maximise model evidence while minimising model complexity, PP systems will use a variety strategies to carry out tasks, from simple heuristics to

---

[6] While Clark (2016, 2017) also discusses PP and EC compatibility, this first analysis is the more sustained and directed work. The subsequent works largely mirror the strategy developed here.

[7] Accuracy refers to the ability to predict incoming sensory information, while complexity refers to the number of parameters used by a model.

complex approximations. This means that in time sensitive tasks, such as the out-fielder problem, reducing prediction errors will often involve the use of models that exploit rolling patterns of perceptual inputs and motor actions.[8] These low-cost models deliver time-sensitive and task-relevant information by taking advantage of various mind-world constancies, such as the optic-acceleration cancellation in the outfielder problem. In this way, PP systems often use problem-solving strategies that are suboptimal but 'good enough' to meet the limited time and processing demands. In other words, RPP naturally accommodates a form of productive laziness. Second, in order to meet task demands in a timely manner, PP systems will also assign high precision to the predictions that underlie actions which enable agents to 'use the world as its own best model'. For example, when performing block-placing tasks, eye saccades allow PP systems to employ minimal-internal-memory strategies. Because co-operating actions with the external environment help to offload computational demands, world-engaging action become suitable for precision-based selection. In this way, RPP also naturally accommodates a form of cognitive scaffolding.

The takeaway is that PP is not only opposed to full-blown (or exclusively) "reconstructive" approaches, such as endorsed by the CPP, but it also accommodates key ideas from EC, such as scaffolding and productive laziness. The RPP offers a systematic way of combining the deep, model based flexibility of PP with the frugal, environmentally-exploitative actions of EC. For Clark, there is a deep compatibility between PP and EC.

2.2 Inferential and Representational Embodiment

A second proposal comes from Hohwy (2018). Hohwy offers a slightly different take on PP-EC compatibility. In line with Clark's proposal, Hohwy also accepts what he sees as an initial tension between elements of PP and EC, such as unconscious perceptual inference and the constitutive role of the body. However, Hohwy further suggests that key notions of EC, such as the flexibility of cognition and the tight coupling between agent and environment, can be folded into the PP scheme.

Hohwy argues that systems engaged in Prediction Error Minimisation over time are properly understood as both inferential and representational. Simplifying slightly, the reasoning is that PEM systems are representational insofar as they need to build up a vast internal model in order to deal with the complex, ever changing environment; while they are inferential insofar as they need to refine their internal models, via approximating a form of Bayesian inference, to minimise long term prediction errors. Hohwy thinks this inferential and representational conception notably reshapes the role and status of the body within PEM systems. This is because bodies are no longer primarily in the business enabling

---

[8] The outfielder problem asks us how a baseball player (an outfielder) is able to catch a fly ball while on the run. It is often used to contrast views which posit internal replica building, such as estimating and tracking the ball's trajectory using rich, internal representations, and more cost efficient methods, such as keeping the image of the ball stationary on their retina in order to keep the flow of sensory information within a certain range.

"attunement" with the world, as some proponent of EC suggest. Rather, they function to aid in the construction and deployment of rich internal representations via the generation of sensory signals.

One supporting observation is that while it is true to say that an agent can minimise prediction error by acting on the world (active inference), this fact nonetheless neglects the formative role played perceptual inference, which allows an agent to learn which actions to choose. It is because an agent can learn that inaction leads to increased prediction error that perceptual inference is necessary. PEM systems learn the error minimising role of action through perceptual inference. What proponents of EC often neglect, Hohwy thinks, is the continuity between perceptual and active inference. This continuity opens up the possibility of explaining two key features of EC: (i) the tight causal coupling of agent and environment and (ii) fast and fluid processing.

In the former's case, Hohwy points out that a set of bodily expectations can be effectively interpreted as a model of a subset of possible states that an organism might occupy. When an organism moves through its environment, the expected states of its model, defined in interoceptive terms, mirror the expected states of the organism, described in environmental, sensory input, or exteroception terms. For example, the sensory organs of fish are more likely to be impinged upon by watery states given that a fish is likely to be found more often than not in water. The expected states that anchor active inference relate to a set of points in an organism's homeostasis. On this basis, Hohwy suggests that perception and cognition cannot be separated from bodily or environmental aspects within PEM systems. Reinterpreted probabilistically, the tight coupling of agent and environment emerges as a natural consequence of the foundational embodiment of PEM systems.

In the latter's case, Hohwy argues that fast and fluid processing can be explained by the gradual build-up of internal representations within PEM systems. He notes that the traditional motivation for thinking about fast and fluid processing is the need to overcome computational bottlenecks in representational systems—affordances, for example, are traditionally thought of as a way to avoid encoding an entire natural scene, which makes action slow and ponderous. PEM systems, in contrast, are said to bypass computational bottlenecks by reconceiving of the role of sensory input. Sensory input does not function to encode natural scenes. Rather, it functions as feedback to the predictive signals generated by internal models. As a result, multi-layered representations are said to build up slowly over time. Fast and fluid processing emerges as the result of fashioning complex expectations about the world using sensory input over time. On Hohwy's view, PEM systems rely on slow and clean learning to facilitate swift and fluid perception and interaction with the world.

So, while Hohwy accepts the initial tension between PP and EC, his proposal departs from Clark (2015) in that it seeks to fold in the key ideas of EC into PP. Writing of the relationship he says, for instance:

> When viewed in this larger context of the free energy principles, promising notions of embodied and embedded condition present themselves. More

research is needed on the extent to which they capture facets of the wide-ranging and heterogeneous 4E body of research. However, for the conception of embodiment and embedding mooted here, an inferential conception is inescapable. (2018, p. 138).

There is room within the PP scheme for EC but only once its key ideas have been given a inferentialist and representational treatment.[9]

3 Three Outstanding Challenges

In the previous section, I outlined two engaging proposals for how to think about the relationship between PP and EC. On the one hand, Clark (2015) envisioned the core elements of EC, such as cognitive scaffolding, fitting comfortably with those of PP, such active inference, in virtue of the embodied-friendly interpretations available to PP. On the other hand, Hohwy (2018) offered a vision in which the key insights of EC, such as agent-environment coupling, could be folded into the PP scheme in light of being given a proper inferential and representational treatment. As we saw, these proposals helped to illuminate a number of interesting conceptual ties. However, what I want to suggest is that, despite their informative character, there are a number of outstanding challenges that yet remain when thinking about PP-EC compatibility.

The first concerns theoretical status. As a number of authors have pointed out, it is often unclear what theoretical status PP and EC are supposed to have within philosophic and scientific theorising. Hohwy (2020), for instance, notes that: "For a framework as broad, detailed, and explanatorily ambitious as PP, there will inevitably be questions about its status in scientific theorising and practice" (p. 220). In a similar vein, Milkowski (2019) writes of EC: "While most surveys, defences, and critiques of embodied cognition proceed by treating it as a neatly delineated claim, such an approach soon becomes problematic due to the inherent plurality of this perspective on cognition." (p. 221). The trouble is that a wide, and largely distinct, array of terms have been applied within discussions of PP and EC, and not always in wholly consistent ways. For example, while some have suggested that PP is a unifying "theory" (Hohwy, 2013; Litwin and Milkowski 2021), others have contended it is a "framework" or "paradigm" for research (van Elk and Aleman 2017; Michel 2022; Sprevak, 2021). Similarly, whereas some have claimed EC is a "thesis" or "hypothesis" (Meteyard et al., 2012; Wilson & Golonka, 2013; Mahon 2015), others have suggested it is a "research programme" or "research tradition" (Shapiro, 2007; Shapiro and Shannon 2021; Milkowski and Nowakowski 2021).

The issue of theoretical status is important, because, as several authors have noted, different theoretical units carry with them different standards of comparison and modes of relation (Michel 2022; Milkowski and Nowakowski 2021). For example, while behaviourism and cognitivism are incompatible at the level of theory,

---

[9] For a discussion of the back-and-forth between Clark and Hohwy on the relationship between PP and EC, see Clark (2017) and Hohwy (2019).

each offering rival attempts to explain specific phenomenon such as language acqui-sition, when viewed at the level of research programme, it is less apparent the views stand in direct tension (Neisser, 1967). As different theoretical units can vary with respect to level of abstractness and function, the character of compatibility can change depending on the type of conceptual unit under discussion. If, for instance, it turns out that PP denotes a "theory-like" structure whereas EC picks out a "para-digm-like" structure, then the question of compatibility may arise in a different form than it would if both denote a "theory-like" structure. While theories can be in direct tension with one another, it is less clear the same is true of theories and larger units of analysis such as paradigms. Because different units of analysis stand in different relations to one another, this complicates any story we want to tell about compatibil-ity. Without first getting clear about the theoretical status of PP and EC, any discus-sion of compatibility may prove premature.

The second issue concerns theoretical commitments. The trouble is that it is often unclear what core commitments PP and EC are supposed to have. For example, in attempting to reconcile PP and EC, Kirchhoff (2018b) articulates four "key" tenets of EC. These include (i) the constitutive thesis, which says that cognitive systems are realised in patterns of sensorimotor activity nonlinearly coupled with the embed-ding environment; (ii) The nonrepresentational thesis, which says that the sensori-motor profile of organisms is sufficient for at least some kinds of cognitive activities; (iii) the cognitive-affective inseparability thesis, the idea that affect, cognition, and sensorimotor contingencies are inseparable; and (iv) the metaplasticity thesis, which says that the entire organism is situated in a plastic network of processes spanning brain, body, and world. Notice, though, that none of these key tenets appear on the lists proposed by Clark (2015) or Hohwy (2018). Clark pointed to "cognitive scaf-folding" and "productive laziness" as key features of EC, whereas Hohwy suggested "agent-environment coupling" and "fast and fluid processing" were central. There is little agreement, even at a very general level, about the core ontological and meth-odological commitments of EC.

A similar point holds in the case of PP. Michel (2022), for example, suggests that the central tenet of PP is that the "mind entertains a probabilistic, hierarchi-cal generative model that aims at anticipating the inflow of sensory information" (p. 6). Sprevak (2021), in contrast, rejects such as characterisation. For Sprevak, PP has little to do with probabilistic inference, generative models or top down effects. While these ideas are used by PP, they do not reflect what is distinctive or unique about PP. Instead, PP should be characterised as a cluster of related claims about the computational, algorithmic and implementational details of cognition (see Sprevak (2021) for details).

The issue of theoretical commitments is important because a lack of consensus could spell trouble for the scope of any account of compatibility. If the commit-ments are spelled out too narrowly, then compatibility might be achieved but only at the cost of sacrificing wider insight. For example, if, as Clark suggests, PP and EC are compatible with respect to cognitive scaffolding and active inference but not, as Kirchhoff suggests, metaplasticity, then Clark's account, while illuminating, would only have limited scope. It would only demonstrate compatibility for specific formulations of EC, such as those pertaining to cognitive scaffolding, but it would

remain silent on how PP relates to other (potential) core commitments of EC, such as metaplasticity. If we are not clear about which commitments are core to EC and PP, then a given account could exclude other relevant commitments. Conversely, if the commitments are defined too broadly, then compatibility could become trivial. For example, if, as is sometimes suggested, EC is simply the commitment to a "crucial role for the body" (Meteyard et al., 2012) and PP is simply the idea that brain is a "prediction machine" (Venter, 2021), then the two views may be compatible but this compatibility is not particularly informative. What we need is a way making sense of the diversity of commitments found within PP and EC, but in a way that does not sacrifice the informativeness of the account of compatibility.

The third issue centres on the notion of compatibility itself. The problem is that while talk of PP and EC compatibility has been relatively commonplace, the sense of compatibility at stake in many discussions has been less clear. For instance, Clark describes compatibility in terms of a type of "fit". The core lessons of EC, such as cognitive scaffolding, are said to fit comfortably together with those of PP, such as active inference, in virtue of the embodied-friendly interpretation available to PP (i.e. the radical versus conservative reading). In contrast, Hohwy describes compatibility in "deflationary" terms. PP-EC compatibility requires rethinking the status of embodied insights along inferential and representational lines.[10] Notice there is a tension looming here. If Hohwy is right, then compatibility is only possible once a strong inferential and representational treatment of PP is adopted; whereas if Clark is right, the opposite is true. What we need is a more precise way of thinking about compatibility. To better understand the relationship between PP and EC, we need to understand the ways in which the two views can be said to be compatible and why. Such an understanding will not only help to clarify differences amongst various proposals, but also point in the direction of how to develop more unified accounts.

So, to sum up, there are three outstanding challenges facing PP-EC compatibility. The first is how to explain the theoretical status of PP and EC; the second is how to specify the theoretical commitments of PP and EC; and the third is how to clarify the sense of compatibility at stake in discussion. To be clear, the point of raising these challenges is not to criticise proposals such as Clark (2015) and Hohwy (2018). Indeed, these proposals, along with others, offer important attempts at clarifying the nature of the relationship between PP and EC. Rather, the point is to draw attention to these outstanding challenges in order to better understand compatibility. As I construe them here, the three issues are different facets of the compatibility issue which have otherwise been overlooked. In resolving the issues, I think we are better placed to understand the compatibility issue more generally. For this reason, I see the current proposal as complementary to Clark and Hohwy's accounts; it expands and deepens the insights already provided.

---

[10] To further speak to the point, Hohwy enumerates three different options for the relationship, all under the heading of "compatibility", but which range from deflationary to reconciliatory approaches.

4 The Problem Solving Model

What I want to suggest is that the three outstanding challenges can be addressed if the issue of PP-EC compatibility is approached using Larry Laudan's (1977, 1981, 1984) "problem solving model" (or PSM for short). The PSM sits within a wider class of models in philosophy of science that attempt to develop historically grounded and conceptually robust accounts of scientific progress and practice; others include "semantic" (Niiniluoto, 1987, 1999) and "epistemic" approaches (Bird, 2007). In particular, the PSM is an instance of the "functionalist-internalist" approach. It is 'functionalist' in that it says that the aim of science is relative to the function it fulfils, and it is 'internalist' insofar as it claims the standards of fulfilment are relative to practitioners' assessments. Unlike epistemic or semantic approaches, the PSM divorces science from "knowledge" and "truth". Instead, the aim of science is to fulfil a certain function.[11] The PSM offers, I think, a number of distinct advantages when used as a meta-theoretic framing for discussion. In operating as a high-level characterisation of scientific theorising and practice, the PSM offers an ideal tool for thinking about the relationship between specific conceptual units within cognitive science such as PP and EC. There are three main components of the model.

The first is the aim of science. As mentioned, according to the PSM, the goal of science is not to deliver "truth" or "knowledge" but, rather, fulfil a certain function: namely, "problem solving". The aim of science is to propose and develop successive theories that solve more "problems" than their predecessors. If, for example, theory A solves more and weightier problems than theory B, then, other things being equal, theory A is preferable to B. According to this picture, scientific progress reflects the successive expansion of the "problem-solving effectiveness" of a theory, or set of theories, within a given domain over time.[12]

The second is the target of analysis. The PSM draws a distinction between two main types of problems within science. The first are "empirical problems". These are various aspects of the natural world which call out for attention, such as why the offspring of plants bear a striking resemblance to their parents. The second are "conceptual problems". These are higher order questions about the conceptual structures used by a theory, such as how behaviourist explanations involved re-descriptions rather than explanations of the intentional aspects of thought. For the PSM, the goal of science is to maximise the scope of empirical problems solved by a theory, and minimise the range of conceptual problems and 'anomalies' that emerge.[13] The problem solving effectiveness of a theory reflects the balance of solved versus unresolved problems. Scientific progress reflects the successive expansion of the

---

[11] For a systematic overview, see Dellsén (2018).

[12] Progress can also occur without the expansion of the solved empirical problems (e.g., when a domain contracts); while regression can occur even if the number of solved problems increases, assuming it leads to more anomalies or conceptual problems accumulating (Laudan, 1977, p. 69).

[13] Anomalies are those problems, either empirical or conceptual, for which a theory is needed but which has not yet been developed.

problem solving effectiveness of a theory or set of theories within a given domain over time.

The third is the unit of analysis. The PSM proposes two main types of conceptual units within science. The first is the more familiar "theory". A theory is set of related doctrines (hypotheses, axioms, and/or principles) which is used to make experimental predictions and provide detailed explanations of natural phenomena—Maxwell's theory of electromagnetism or Einstein's theory of photoelectric effect offer classic examples within physics. A theory articulates a concrete ontology and number of specific and testable claims about the world. A successful theory is one which adequately addresses specific empirical and conceptual problems within a given domain.

The second is a broader unit, what Laudan calls a "research tradition". A research tradition functions to guide, inspire, constrain, and rationalise theory construction and development. It operates at a more abstract level than a theory, specifying what the world is made of and how it should be studied. It is composed of a set of component theories and certain metaphysical and methodological assumptions—Newton's theory of gravity or Hutton's uniformitarianist theory of geologic time offer classic examples. A successful research tradition is one which leads, via its component theories, to the solution of an increasing range of empirical and conceptual problems. The assessment of a research tradition is bound up with the problem-solving effectiveness of the entire set of theories.

A research tradition is distinguished by three main characteristics (Laudan, 1977, p. 78–9). First, it is associated with a number of specific theories. Skinner's theory of language acquisition, for example, offers an exemplifying instance of the behaviourist research tradition. Second, a research tradition possesses a number of ontological and methodological commitments. As mentioned, these commitments specify the fundamental entities and the appropriate methods for inquiry.[14] If the research tradition is behaviourism, for instance, then the commitments include studying directly observable physical and physiological entities. Third, a research tradition has a long and detailed history, i.e. it usually takes on several, sometimes contradictory, formulations over time. For example, earlier forms of Watsonian behaviourism departed significantly from later formulations by Skinner with respect to the source of conditioning, e.g., pre- versus post-stimulus reinforcement. In short, theories serve to address specific empirical and conceptual problems within a given domain, while research traditions serve as broader units of scientific change, guiding theory construction and establishing the continuity of science.

While there is more that can be said about the PSM, there are three important takeaways for the moment: (i) the aim of science is to maximise problem solving effectiveness; (ii) science advances when the range of empirical and conceptual problems solved increases over time; and (iii) theories and research traditions

---

[14] In practice, the ontological and methodological commitments are often intertwined. This is because the appropriate methods of inquiry are generally compatible with the objects of inquiry.

constitute the main units of analysis within science.[15] Moreover, as an outright defense of the PSM is beyond the scope of the current discussion, I assume in what follows that the case for the PSM is further bolstered by the work it does in resolving the outstanding challenges.[16]

5 A Model Solution

With the PSM outlined, I want to return now to the three outstanding challenges.

First, notice that the PSM offers a clear way of thinking about the theoretical status of PP and EC. Namely, PP and EC emerge as "research traditions" on the PSM. To see why, recall the three main characteristics of research traditions.

First, there are a number of specific theories associated with PP and EC. PP theories, for instance, have been offered for everything from body perception (Apps and Tsakiris 2014) and attention (Hohwy, 2012) to psychiatric disorders (Friston et al., 2014) and psychedelics (Deane, 2021). Similarly, detailed embodied accounts have been developed for a range of topics, including language processing (Lakoff & Johnson, 1999), visual phenomenology (Noë, 2009), and music cognition (Kersten & Wilson, 2016; Leman, 2007), to name a few. These detailed accounts particularise the various ontological and methodological assumptions articulated by the wider tradition, including those about sensorimotor knowledge or prediction error minimisation. Both PP and EC possess a number of exemplifying or (partially) constitutive instances.

Second, PP and EC both possess a number of ontological and methodological assumptions. For example, as mentioned, PP regularly employs talk of generative models, efficient neural coding, prediction error minimisation, and top-down effects.[17] In so doing, it commits itself to a particular hierarchical predictive structure for cognitive/neural systems. It proscribes a particular set of processes and entities for study. What is more, it employs specific forms of computational modelling—many PP theories, for example, model cognitive systems as a form Bayesian or probabilistic inference. These models treat cognitive systems as approximating a form of Bayes' rule in order to minimise prediction error. In this way, it also proscribes general methods or principles for inquiry. Similarly, as Milkowski and Nowakowski (2021) point out, EC exhibits a number of metaphysical and methodological commitments within its subtradition. For example, within sensorimotor

---

[15] Also, as a point of clarification, while I only address "theories" and "research traditions" in what follows, this does not mean that these are the only conceptual units within scientific discourse and practice. I readily acknowledge, for example, that there is an important place for talk of models within science. Nonetheless, I think this omission is justified at present because the literature is mostly in agreement that models operate below the level of theories, and so they do not affect the main point I want to make in what follows. For this reason, I use the terms largely as interchangeable.

[16] For critical discussion of the model, see Bird (2007) or Rowbottom (2010).

[17] In PP, a generative model is a statistical model of how data are generated, which consists of prior distributions over the environmental causes of agents' sensory data and generative distributions (or likelihoods) of agents' sensory data given their environmental causes.

theories, such as O'Regan and Noë (2001) and Thelen et al. (2001), cognitive processes are said to involve complex sensorimotor contingencies, rather than veridical representations. Moreover, these theories often employ dynamical systems theory as a way of modelling the dynamic, world-involving actions of cognitive agents (Chemero, 2009). In this way, EC also offers methods of inquiry for its entities.

To clarify, while individual theories also exhibit ontological and methodological commitments, it is the diversity and generality of commitments found within PP and EC that speak in favour of their interpretation as research traditions. There are a cluster of ideas, for example, which, while not found in every instance, are present within a number of PP models, such as prediction error minimisation, generative models, and probabilistic inference. The easiest way to accommodate this range of concepts and principles is to conceptualise them not only as particular assumptions within specific theories, but also organising or guiding principles for research.[18]

Finally, PP and EC have both undergone a series of transformations over time, particularly as new theories have developed and older ones have been modified or dropped. PP, for example, charts its evolution from the early inferentialist work of Helmholtz and top-down approaches of Gestalt psychology to more recent developments in Bayesianism (Hohwy, 2013, 2016) and active inference (Friston et al., 2018; Constant et al., 2020); while EC, with its early roots in cybernetics and ecological psychology, has evolved to incorporate various elements of connectionism (Clark, 1997), developmental psychology (Thelen, 1995), and dynamical systems theory (Beer, 2000). Over time, PP and EC have evolved to incorporate new ontological and methodological commitments.

So, to sum up, in virtue of exemplifying its key characteristics, PP and EC should be thought of as research traditions. Both views operate as general units within science that (i) possess a number of exemplifying instances, (ii) exhibit various metaphysical and methodological assumptions, and (iii) have distinct developmental histories. The PSM articulates a clear answer to the theoretical status issue.[19]

I am also not alone in this position. Michel (2022), for instance, suggests that PP is a "cognitive computational *paradigm* that is only just emerging and still under construction" (original emphasis) (p.5); while Miłkowski and Nowakowski (2021) suggest that, "EC should be understood as composed of multiple, and sometimes quite extensive, component research subtraditions" (p. S71). However, while landing on the same core idea these accounts differ from the current one with respect to their target of analysis. Michel (2022), for instance, is concerned with articulating the theoretical status of PP in order to better understand the language of thought hypothesis, while Miłkowski and Nowakowski (2021) are concerned with the status of EC

---

[18] A similar point holds of EC. See Miłkowski and Nowakowski (2021) for discussion.

[19] One interesting point to note is that when a specific theory is abstracted from its historical context this can produce ambiguous clues as to its research tradition (or traditions) (Laudan, 1977, p. 99). This can lead some scientists and philosophers to think that certain theories are assessed independently of their parent tradition. This seems to be what has happened to some degree in the case of PP. In a rush to emphasise its "novelty" and "promise", discussions have tended to downplay PP's theoretical continuity with other traditions. This is understandable to extent, but it has contributed, I think, to some of the confusion surrounding PP's theoretical status and its commitments.

in order to further understand representational unification. In contrast, the current proposal is interested in theoretical status in order to better elucidate the nature of the relationship between PP and EC.

Next, notice that the PSM helps to explain why the core commitments of PP and EC have remained so elusive. First, recall that a research tradition specifies, in a very general way, the basic types of fundamental entities that exist within a domain, as well as methods for inquiry; whereas a theory articulates a specific ontology and a number of testable claims about nature. Second, recall that one of the primary functions of a research tradition is to guide and constrain theory construction. As Laudan (1977, p. 84) notes, these facts together imply that for any given theory there will be a variety of ways it might implement the basic commitments of its parent tradition. For example, while the Newtonian research tradition claims that all non-rectilinear motions should be treated as cases of centrally directed forces, this commitment does not entail how a specific theory should explain the motion of a compass when it is near a current-carrying wire. To develop a Newtonian theory of that particular phenomenon, a researcher would need to go beyond the general commitments of the parent tradition.

One consequence of this is that one cannot read off the commitments of a research tradition from its individual theories, nor can the individual commitments of a theory be read off its parent tradition. This explains why it has proven so difficult to clarify the commitments of PP and EC. If one were to survey embodied theories of semantics, for example, one would be forgiven for thinking that a notion of "representation" is common to EC; talk of the sensory and motor information in cognitive representations is an constant theme in discussions of embodied semantics (see, e.g., Barslou, 2008; Lakoff, 2012; Dove, 2022). However, as is well known, a number of distinct EC theories, such as Schoner and Thelen (2006) or Noë (2009), explicitly eschew talk of representations in favour of explanations involving sensorimotor contingencies and action-dynamics. Similarly, a survey of PP reveals a number of theories which invoke talk of generative models, such as Corlett et al. (2019), but others still, such as Kiverstein et al. (2019) that do not. Because different theories have different ways of implementing a research tradition's general commitments, attempting to infer the core commitments from a survey of associated theories is not only unlikely to succeed but also potentially misleading.

Another complicating factor is that while a research tradition possesses a characteristic set of ontological and methodological commitments at any one time, the central elements will continue to change over time (c.f. Lakatos, 1970). For example, what was initially taken to be an ineliminable part of Newtonian physics in the seventeenth century (absolute space and time) was no longer regarded as central by the mid-nineteenth century (Laudan, 1977, p. 99). There is relative but not complete continuity between the central elements of a research tradition.

This is also true in the case of PP and EC. For example, as Sprevak (2021) points out, there are a number of ideas which are often invoked in the context of PP but which are also shared by a variety of alternative computational approaches to cognition. These include the idea that (i) the brain employs an efficient coding scheme, (ii) cognition contains many top-down, expectation-driven effects, (iii) cognition involves minimising prediction error, (iv) cognition is a form of probabilistic

inference, and (v) cognition employs generative models. As Sprevak sees matters: "If one wishes to know what is novel with predictive coding [PP], then these ideas, whatever their value, can function as potential distractors." As PP has evolved, conceptual relations have become untangled and various commitments initially seen as central to the view have moved to the periphery. Because there is no complete preservation of a research tradition's core elements, a snapshot of a tradition's theories at any one time can give a misleading impression about its core commitments. This is another potential source of confusion.

Of course, one may have noticed that I have avoided explicitly stating what the theoretical commitments of PP and EC are. There are two reasons for this. The first is that while PP is a research tradition in the sense proposed by the PSM, it is still relatively early in its development. This means that it is difficult to specify its core commitments with any precision. As Laudan (1977, p. 92) points out, a research tradition's commitments are only explicable retrospectively, and so it may prove premature to try to identify PP's commitments in its current state.[20] The second is that identifying the *specific* commitments of a tradition, rather than the presence of commitments more broadly, requires a detailed survey of the tradition's historical development, and such a detailed analysis of EC would go well beyond the scope of current paper.

What has been offered, though, are the tools by which to identify the core commitments of PP and EC. To arrive at the core commitments, researchers must (i) avoid making inferences solely on the basis of a research tradition's associated theories or models and (ii) provide a rich, detailed study of PP and EC's developmental history.[21] While this is a slightly more modest result than some might desire, it is nonetheless a relevant contribution given the general state of uncertainty surrounding PP and EC's commitments.

Finally, notice that the PSM offers a more precise way of thinking about the concept of compatibility. As we saw, the PSM articulates two main units of analysis: theories and research traditions. This distinction proved helpful in sorting out the theoretical status of PP and EC and diagnosing the confusion surrounding PP and EC's theoretical commitments. But, in addition to this, the PSM helps to spell out how the two basic units of science can be said to relate.

According to the PSM, research traditions stand in two basic relations to one another. The first is *consistency*; the second *integration*. In the former's case, two or more research traditions are consistent, and therefore compatible, if their metaphysical and methodological commitments do not contradict. In the latter's case, two or more research traditions can be integrated if elements of one research tradition can be blended or combined with those of another. Integration comes in two forms. The first, what Laudan (1977, p. 103) calls "grafting", occurs when two or more research traditions are integrated without the core elements of one tradition being undermined by those of another. The second, what Laudan (1977, p. 104)

---

[20] This is point is echoed by Michel (2022, p. 5).
[21] For an instructive start in the case of EC, see Milkowski and Nowakowski (2021).

calls "repudiating", occurs when one or more elements of a research tradition are dropped in favour of another's.

What is interesting about this taxonomy is that it offers a close fit with existing proposals. For instance, according to Clark's proposal, the radical reading of PP offers a systematic way of the combining the deep, model based flexibility of PP with the frugal, environmentally-exploitative actions of EC. As Clark (2015, p. 24) makes the point: "The worry that predictive processing organisation might over-emphasise computationally expensive, representation heavy strategies over quicker, dirtier, more 'embodied' ones is thus fully and satisfyingly resolved. The ever-active predictive brain stands revealed as a lazy brain – a brain vigilant for any opportunity to do less while achieving more." Clark's proposal focuses on showing that core elements of EC, such as cognitive scaffolding, can be combined with central elements of PP, such precision estimation, without the former undermining the latter. In this way, it offers a representative example of integration via grafting. PP and EC can be amalgamated or combined once the proper interpretation of PP is settled (i.e. the radical versus conservative reading of PP).

Hohwy's proposal, on the other hand, attempts to show that key the insights and ideas of EC, such as tight agent-environment coupling and fast and frugal processing, can be accommodated within the PP scheme but only once they have been reinterpreted probabilistically. As Hohwy (2018) expresses the point: "Perhaps the basic sentiment could be summed up in the strong intuition that embodied action is not inference, and yet the body and its actions are crucial to gain any kind of understanding of perception and cognition. PEM can, however, easily cast action as a kind of inference—as active inference." In adopting a deflationary approach to compatibility, Hohwy's proposal offers a clear example of integration by repudiation. PP and EC can be combined, but only if the non-inferential and non-representational elements of EC are dropped.

Finally, consider Kirchhoff's (2018b) account. As alluded to, Kirchhoff argues that PP and EC are compatible insofar as core elements of PP do not contradict the four key "tenets" of EC. As mentioned, these include: (i) the constitutive thesis, (ii) the non-representational thesis, (iii) the cognitive-affectivity inseparability thesis, and (iv) the metaplasticity thesis. For reasons of space, I cannot summarise all the argumentation in detail, but the essential line, similar to that of Clark, is that there are interpretations available for key PP notions, such as inference and generative models, which avoid any inconsistency with the key tenets of EC. For instance, writing about the inseparability thesis, Kirchhoff (2018b) notes: "…on the assumption that affectivity is an essential component of perception, and if sense-making is inherently affective, then for organisms to enact their world in prediction error minimization is also for organisms to enact it affectively". For Kirchhoff, there is no incompatibility between PP and EC, because the core tenets of each do not contradict. Here, again, we have a nice example of the PSM articulating the form of compatibly at play: namely, consistency.

In short, the PSM helps to flesh out the different senses of compatibility at stake in discussion. Differences amongst proposals emerge as a function of differences in the form of compatibility adopted, e.g., consistency versus integration. The PSM not

only offers a descriptively adequate fit with existing proposals, but it also offers a further way of nuancing the discussion of PP-EC compatibility.

So, taking stock, the PSM offers the resources to deal with each of the three outstanding challenges. First, it resolves the theoretical status issue in virtue of revealing PP and EC as "research traditions"; second, it resolves the theoretical commitments issue by diagnosing the sources of existing ambiguity; and third, it resolves the ambiguity surrounding the concept of compatibility issue by fleshing out the different senses of compatibility at stake within discussion.

To be clear, though, in having resolved the three outstanding issues, I have not thereby settled the compatibility issue. As mentioned, the three outstanding issues are different facets of the compatibility issue, ones which have otherwise largely been overlooked. It still remains to be seen which proposal ultimately proves correct, e.g., Clark (2015), Hohwy (2018), Kirchhoff (2018b). What I have provided, though, are additional tools for developing such proposals. As we have seen, framing discussion through the PSM offers a productive way of thinking about not only how to identify PP and EC's theoretical commitments but also how compatibility can be achieved.

6 Further Implications

To conclude, it will be worth fleshing out two implications of the PSM for PP and EC more generally.

One consequence is that the PSM provides a productive way of understanding the often varied character of PP and EC discussions. Shapiro (2007), for instance, suggests that EC is a research programme, but spends considerable time describing specific examples of EC theories, such as Glenberg and Robertson's (2000) work on the symbol grounding problem or Thelen and Smith's (1994) work on infant motor development. Similarly, Clark (2016) maintains that PP is a framework, but his discussion provides a number of detailed reviews of specific PP theories, such as those about vision and attention. One consistent theme within discussions of PP and EC is a shift between low- and high-level descriptions. Notice that these shifts are quite understandable when viewed through the lens of the PSM. This is because they reflect a natural move between the *theory* and *research tradition* senses of PP and EC. On the one hand, talk of detailed and testable accounts reflects the theory sense of PP and EC; while the broader, more ambitious talk reflects the research tradition sense.

A second interesting implication is that the PSM helps to address several recent critiques of PP and EC. Litwin and Miłkowski (2021), for instance, worry that PP not only does not adequately justify its fundamental tents but it also involves models which are mutually inconsistent with its fundamental tenets. They write, for instance: "[i]nstead of developing 'vertically,' or simply going deeper into fundamentals of the theory to increase its theoretical virtues, a plethora of proto-models and theories—frequently mutually exclusive and inconsistent with basic PP tenets— is being formulated in liberally interpreted PP terms."

However, such a worry is misplaced once the PSM has been adopted. First, notice that it is not the place of specific theories to rationalise or justify their core assumptions; theories are not "self-authenticating". Rather, as Laudan (1977, p. 79) points out, this is one of the functions of research traditions. It is the research tradition which specifics, in broad terms, the objects and methods of inquiry. As we saw, since PP is best understood as a research tradition, Litwin and Miłkowski are off the mark in criticising it for having unclear fundamentals – recall also the difficulties in identifying those fundamentals given PP's relatively early stage of development. Michel (2022) makes a similar point when he writes: "[I]f we characterise PP as a *paradigm* then such criticisms miss the point. What might deserve those criticisms are, of course, specific theories of specific phenomena that make use of the core concepts and principles of PP" (p. 5). Interpreted along the lines of the PSM, Litwin and Miłkowski's critique is more an invitation to reflect on the conceptual well-foundedness of PP's metaphysical and methodological commitments, rather than an indictment of those commitments.[22]

Second, notice that even if specific PP theories or models come into conflict with the wider tenets of the research tradition, then this would still be okay according to the PSM. This is because, as individual theories develop and change to tackle an increasing range of conceptual and empirical problems, inconsistencies may arise with other theories of the tradition, or even core tenets of the tradition itself. This is a natural part of the evolution of research tradition; as behaviourist theories evolved in response to criticism, for example, some began to incorporate talk of latent, internal responses. While some theories may conflict with core commitments of the research tradition, this does not thereby undermine their epistemic value. What matters for the evaluation of a research tradition is its problem solving effectiveness.

A related criticism has been levelled against EC. Goldinger et al. (2016), for instance, argue that EC is theoretically vacuous in that it fails to predict or explain numerous classic phenomena in cognitive science, such as the word frequency effect. They write, for instance: "If one adopts the stance that cognition is fundamentally rooted in bodily states, a vast array of data are immediately beyond hope of theoretical explanation" (2016, p. 974).

But, again, such a criticism only makes sense if it targets the theory rather than research tradition sense of EC. While individual EC theories, such as those about word frequency effect, are empirically falsifiable, in virtue of making specific experimental predictions, research traditions, by their nature, are not testable. For example, as Miłkowski and Nowakowski (2021) point out, computationalism, like EC, is not a single theory but, rather, a varied tradition, one methodologically committed to computational modelling and ontologically committed to claims about physical

---

[22] Another possibility, one which I do not have space to explore, is that Litwin and Miłkowski (2021) are raising the prospect of PP being a "nonstandard" research tradition (Laudan, 1977, pp. 105–106). Nonstandard research traditions are those which, while missing one or more core features, possess a genuine coherence – for example, early psychometrics, while not possessing a substantive ontology, was held together by a more general conviction that mental phenomena could be represented mathematically. If PP was a nonstandard research tradition, this would also help to explain why its theoretical assumptions remain underdeveloped.

computation in cognitive systems. It does not make detailed predictions about specific phenomenon such as the word frequency effect. Rather, it specifies experimental procedures and modes of inquiry for investigating cognitive phenomenon. It provides a guide to experiment, but is itself not directly testable. If EC is empirically unfalsifiable, then so too is computationalism.

7 Conclusions

The goal of this paper has been to further clarify the nature of PP-EC compatibility. I sought to do so by addressing three outstanding challenges. The first was how to explain the theoretical status of PP and EC; the second was how to specify the theoretical commitments of PP and EC; and the third was how to clarify the sense of compatibility at stake in discussion. In response to these challenges, I introduced the PSM. After outlining its key components, I argued that the PSM offered a clear route to addressing all three challenges. First, it addressed the theoretical status issue in virtue of revealing PP and EC to be "research traditions"; second, it addressed the theoretical commitments issue by diagnosing the sources of existing ambiguity; and third, it addressed the compatibility issue by fleshing out the various senses of compatibility at stake within discussion. I also outlined further implications of adopting the PSM. These included explaining the varied character of PP and EC discussions and responding to several recent criticisms. These further implications are important as they point in direction of relevant future work. As I see it, the major contribution of this paper lies not only in the specific answers it offers, but also in the structure it provides to discussion. In reframing compatibility in terms of the PSM, I hope to set have discussion of PP and EC on a more constructive and well-delineated path going forward.

References

Anderson, M. L. (2017). Of Bayes and Bullets: An embodied, situated, targeting-based account of predictive processing. In T. Metzinger & W. Wiese (Eds.), *Philosophy and Predictive Processing: 4.* MIND Group.

Apps, M. A. J., & Tsakiris, M. (2014). The free-energy self: A predictive coding account of self-recogni-tion. *Neuroscience & Biobehavioral Reviews, 41*, 85–97.

Ballard, D., Hayhoe, M., Pook, P., & Rao, R. (1997). Deictic codes for the embodiment of cognition. *Behavioral and Brain Sciences, 20*, 4.

Barsalou, L. W. (2008). Grounded cognition. *Annual Review of Psychology, 59*, 617–645. https://doi.org/10.1146/annurev.psych.59.103006.093639

Beer, R. (2000). Dynamical approaches to cognitive science. *Trends in Cognitive Science, 4*(3), 91–99.

Bird, A. (2007). What is scientific progress? *Noûs, 41*(1), 64–89. https://doi.org/10.1111/j.1468-0068.2007.00638.x

Bissell, D. A., Ziadni, M. S., & Sturgeon, J. A. (2018). Perceived injustice in chronic pain: An examina-tion through the lens of predictive processing. *Pain Management, 8*(2), 129–138. https://doi.org/10.2217/pmt-2017-0051

Bruineberg, J., Rietveld, E., & Kiverstein, J. (2018). The anticipating brain is not a scientist: The free-energy principle from an ecological-enactive perspective. *Synthese, 195*, 2417–2444. https://doi.org/10.1007/s11229-016-1239-1

Chemero, A. (2009). *Radical embodied cognitive science*. The MIT Press.

Clark, A. (1997). *Being there: Putting Brain*. MIT Press.

Clark, A. (2008). *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford Univer-sity Press.

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences, 36*(3), 181–204.

Clark, A. (2015). Radical predictive processing. *The Southern Journal of Philosophy, 53*(S1), 3–27. https://doi.org/10.1111/sjp.12120

Clark, A. (2016). *Surfing uncertainty: Prediction, action, and the embodied mind*. Oxford University Press.

Clark, A. (2017). Busting out: Predictive brains, embodied minds, and the puzzle of the evidentiary veil. *Noûs, 51*(4), 727–753.

Constant, A., Clark, A., Kirchhoff, M. D., & Friston, K. J. (2020). Extended active inference: construct-ing predictive cognition beyond skulls. *Mind and Language*. https://doi.org/10.1111/mila.1233

Corlett, P. R., Horga, G., Fletcher, P. C., Alderson-Day, B., Schmack, K., & Powers, A. R. (2019). Hal-lucinations and strong priors. *Trends in Cognitive Sciences, 23*(2), 114–127.

Deane, G. (2021). Consciousness in active inference: Deep self-models, other minds, and the challenge of psychedelic-induced ego-dissolution. *Neuroscience of Consciousness, 2021*(2), niab024.

Dellsén, F. (2018). Scientific progress: Four accounts. *Philosophy Ompass, 13*(11), e12525. https://doi.org/10.1111/phc3.12525

Dove, G. (2022). *Abstract Concepts and the Embodied Mind: Rethinking Grounded Cognition*. Oxford University Press.

Friston, K. J. (2018). Does predictive coding have a future? *Nature Neuroscience, 21*(8), 1019–1021. https://doi.org/10.1038/s41593-018-0200-

Friston, K. J., Stephan, K. E., Montague, R., & Dolan, R. J. (2014). Computational psychiatry: The brain as a phantastic organ. *The Lancet Psychiatry, 1*(2), 148–158. https://doi.org/10.1016/S2215-0366(14)70275-5

Gallese, V., & Lakoff, G. (2005). The brain's concepts: The role of the sensory-motor system in concep-tual knowledge. *Cognitive Neuropsychology, 22*(3/4), 455–479.

Glenberg, A. M., & Robertson, D. (2000). Symbol grounding and meaning: A comparison of high-dimensional and embodied theories of meaning. *Journal of Memory and Language, 43*, 379–401.

Goldinger, S. D., Papesh, M. H., Barnhart, A. S., Hansen, W. A., & Hout, M. C. (2016). The poverty of embodied cognition. *Psychonomic Bulletin & Review, 23*(4), 959–978. https://doi.org/10.3758/s13423-015-0860-1

Hohwy, J. (2012). Attention and conscious perception in the hypothesis testing brain. *Frontiers in Psy-chology, 3*, 96. https://doi.org/10.3389/fpsyg.2012.00096

Hohwy, J. (2013). *The predictive mind*. Oxford University Press.

Hohwy, J. (2016). The self-evidencing brain. *Nous, 50*(2), 259–285.

Hohwy, J. (2018). The Predictive Processing Hypothesis. In A. Newen, L. De Bruin, & S. Gallagher (Eds.), *The Oxford Handbook of 4E Cognition* (pp. 129–145). Oxford University Press.

Hohwy, J. (2019). Quick 'n' Lean or Slow and Rich?: Andy Clark on Predictive Processing and Embod-ied Cognition. In M. Colombo, E. Irvine, & M. Stapleton (Eds.), *Andy Clark and His Critics* (pp. 191–205). Oxford University Pres. https://doi.org/10.1093/oso/9780190662813.003.0015

Hohwy, J. (2020). New directions in predictive processing. *Mind & Language, 35*, 209–223.

Hutto, D. (2018). Getting into predictive processing's great guessing game: Bootstrap heaven or hell? *Synthese, 195*, 2445–2458. https://doi.org/10.1007/s11229-017-1385-0

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

Hutto, D., & Myin, E. (2017). *Evolving enactivism: Basic minds meet content*. MIT Press.

Kersten, L. (2017). Extended music cognition. Philosophy. *Psychology, 30*(8), 1078–1103.

Kersten, L. (2022). A new mark of the cognitive? Predictive processing and extended cognition. *Synthese, 200*(4), 1–25.

Kersten, L., & Wilson, R. (2016). The sound of music: Externalist style. *American Philosophical Quarterly, 53*(2), 139–154.

Kirchhoff, M. (2018a). Predictive brains and embodied, enactive cognition: An introduction to the special issue. *Synthese, 195*, 2355–2366.

Kirchhoff, M. (2018b). The Body in Action: Predictive Processing and the Embodiment Thesis. In A. Newen, L. De Bruin, & S. Gallagher (Eds.), *The Oxford Handbook of 4E Cognition* (pp. 243–260). Oxford University Press.

Kirchhoff, M., & Kiverstein, J. (2019). *Extended Consciousness and Predictive Processing: A Third Wave View*. Routledge.

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

Lakatos, I. (1970). Falsification and the methodology of scientific research programmes. In I. Lakatos & A. Musgrave (Eds.), *Criticism and the growth of knowledge: Proceedings of the international colloquium in the philosophy of science*. Cambridge University Press.

Lakoff, G. (2012). Explaining embodied cognition results. *Topics in Cognitive Science, 4*(4), 773–785. https://doi.org/10.1111/j.1756-8765.2012.01222.x

Lakoff, G., & Johnson, M. (1999). *Philosophy in the flesh: The embodied mind and its challenge to western thought*. Basic Books.

Laudan, L. (1977). *Progress and its problem: Towards a theory of scientific growth*. University of California Press.

Laudan, L. (1981). A problem-solving approach to scientific progress. In I. Hacking (Ed.), *Scientific revolutions* (pp. 144–155). Oxford University Press.

Laudan, L. (1984). *Science and Values*. University of California Press.

Leman, M. (2007). *Embodied Music Cognition and Mediation Technology*. MIT Press.

Litwin, P., & Miłkowski, M. (2021). Unification by fiat: Arrested development of predictive processing. *Cognitive Science*. https://doi.org/10.1111/cogs.12867

Mahon, B. Z. (2015). The burden of embodied cognition. *Canadian Journal of Experimental Psychology, 69*(2), 172–178. https://doi.org/10.1037/cep0000060

Meteyard, L., Rodriguez, S., Bahrami, B., & Vigliocco, G. (2012). Coming of age: A review of embodiment and the neuroscience of semantics. *Cortex, 48*(7), 788–804.

Michel, C. (2022). Scaling up predictive processing to language with construction grammar. *Philosophical Psychology*. https://doi.org/10.1080/09515089.2022.2050198

Miłkowski, M. (2019). Fallible heuristics and evaluation of research traditions. The case of embodied cognition. *Ruch Filozoficzny, 75*(2), 223–236. https://doi.org/10.12775/rf.2019.031

Miłkowski, M., & Nowakowski, P. (2021). Representational unification in cognitive science: Is embodied cognition a unifying perspective? *Synthese, 199*(Suppl 1), 67–88.

Miller, M., & Clark, A. (2018). Happily entangled: Prediction, emotion, and the embodied mind. *Synthese, 195*(6), 2559–2575. https://doi.org/10.1007/s11229-017-1399-7

Neisser, U. (1967). *Cognitive Psychology*. Appleton-Century-Crofts.

Niiniluoto, I. (1987). *Truthlikeness*. Reidel.

Niiniluoto, I. (1999). *Critical Scientific Realism*. Oxford University Press.

Noë, A. (2004). *Action in Perception*. MIT Press.

Noë, A. (2009). *Out of Our Heads: Why you are not your Brain, and other lessons from the Biology of Consciousness*. Hill and Wang.

O'Regan, J. K., & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *The Behavioral and Brain Sciences, 24*(5), 939–973. https://doi.org/10.1017/S0140525X01000115

Rowbottom, D. (2010). What scientific progress is not: Against bird's epistemic view. *International Studies in the Philosophy of Science, 24*(3), 241–255.

Schöner, G., & Thelen, E. (2006). Using dynamic field theory to rethink infant habituation. *Psychological Review, 113*(2), 273–299. https://doi.org/10.1037/0033-295X.113.2.273

Seth, A. K., Suzuki, K., & Critchley, H. D. (2012). An interoceptive predictive coding model of conscious presence. *Frontiers in Psychology, 2*, 395. https://doi.org/10.3389/fpsyg.2011.00395

Shapiro, L., and Spaulding, S. (2021). Embodied Cognition. The Stanford Encyclopaedia of Philosophy. Edward N. Zalta (ed.). https://plato.stanford.edu/archives/win2021/entries/embodied-cognition/. Accessed 11 Oct 2022.

Shapiro, L. A. (2007). The embodied cognition research programme. *Philosophy Compass, 2*(2), 338–346. https://doi.org/10.1111/j.1747-9991.2007.00064.x

Shapiro, L. (2011). *Embodied cognition*. Routledge.

Sprevak, M. (2021). Predictive coding II: The computational level. TBC. [Online access].

Thelen, E. (1995). Motor development: A new Synthesis. *American Psychologist, 50*, 79–95.

Thelen, E., & Smith, L. (1994). *A dynamical systems approach to the development of perception and action*. Cambridge, MA: MIT Press.

Thelen, E., Schöner, G., Scheier, C., & Smith, L. B. (2001). The dynamics of embodiment: A field theory of infant perseverative reaching. *Behavioral and Brain Sciences, 24*(1), 1–34. https://doi.org/10.1017/S0140525X01003910

van Elk, M., & Aleman, A. (2017). Brain mechanisms in religion and spirituality: An integrative predictive processing framework. *Neuroscience & Biobehavioral Reviews, 73*, 359–378. https://doi.org/10.1016/j.neubiorev.2016.12.031

Varela, F., Thompson, E. E., & Rosch. (1991). *The Embodied Mind*. MIT Press.

Venter, E. (2021). Toward an embodied, embedded predictive processing account. *Frontiers in Psychology, 12*, 543076. https://doi.org/10.3389/fpsyg.2021.543076

Ward, D., Silverman, D., & Villalobos, M. (2017). Introduction: The varieties of enactivism. *Topoi, 36*(3), 365–375. https://doi.org/10.1007/s11245-017-9484-6

Williams, D. (2019). Hierarchical minds and the perception/cognition distinction. *Inquiry*. https://doi.org/10.1080/0020174X.2019.1610045

Wilson, R. A. (2004). *Boundaries of the Mind: The Individual in the Fragile Science*. Cambridge University Press.

Wilson, R. A., & Clark, A. (2009). How to situate cognition: Letting nature take its course. In P. Robbins & M. Aydede (Eds.), *The Cambridge handbook of situated cognition* (pp. 55–77). Cambridge University Press.

Wilson, A. D., & Golonka, S. (2013). Embodied cognition is not what you think it is. *Frontiers in Psychology*. https://doi.org/10.3389/fpsyg.2013.0005

Wilson, R. A., & Foglia, L. (2017). Embodied cognition. In E. N. Zalta (Ed.), *Stanford encyclopedia of philosophy (Spring 2017)*. Metaphysics Research Lab, Stanford University.