
Extending the Explanandum for Predictive Processing

A Commentary on Andy Clark

Michael Madary

In this commentary, I suggest that the predictive processing framework (PP) might be applicable to areas beyond those identified by Clark. In particular, PP may be relevant for our understanding of perceptual content, consciousness, and for applied cognitive neuroscience. My main claim for each area is as follows:

- 1) PP urges an organism-relative conception of perceptual content.
- 2) Historical *a priori* accounts of the structure of perceptual experience converge with results from PP.
- 3) There are a number of areas in which PP can find important practical applications, including education, public policy, and social interaction.

Keywords

Anticipation | Applied cognitive neuroscience | Consciousness | Perception | Perceptual content | Phenomenology | Predictive processing

Commentator

Michael Madary

madary@uni-mainz.de

Johannes Gutenberg-Universität
Mainz, Germany

Target Author

Andy Clark

Andy.Clark@ed.ac.uk

University of Edinburgh
Edinburgh, United Kingdom

Editors

Thomas Metzinger

metzinger@uni-mainz.de

Johannes Gutenberg-Universität
Mainz, Germany

Jennifer M. Windt

jennifer.windt@monash.edu

Monash University
Melbourne, Australia

1 Introduction

An understandable reaction to the predictive processing framework (PP) is to think that it is too ambitious ([Hohwy this collection](#)). My suggestion in this commentary is the opposite. I will argue that PP can be fruitfully applied to areas of inquiry that have so far received little, if any, attention from the proponents of PP. Perhaps we can extend the explanandum even further than Andy Clark has recommended.

There is a certain rhetorical danger to the position I am urging. One should not oversell

one's case. I hope to avoid this danger by being clear upfront that my goal is not to convince the skeptic of the attraction of PP. I cannot improve on Clark (and others, see below) in that regard. Instead, I investigate the following question: if some version of PP (again, see below) is true, then what are the larger implications for human self-understanding? My answer to this question covers three topics. First I will engage with Clark's discussion of perceptual processing from sections 1 and 2.1 of his article. There I

will sketch how PP's reversal of the traditional model of perceptual processing may have significant implications for the way in which we understand perceptual content, which is a core issue in the philosophy of psychology. In the [second](#) section I will turn to another area of philosophical concern: consciousness. Historically, consciousness research has had a rocky relationship with the sciences of the mind. I hope to point towards the possibility of a rapprochement. In the final section of the commentary, I will quickly touch on some practical matters. If PP is true, then there are important consequences for the way in which we approach topics in education, public policy, and social interaction.

My goal is to indicate possible areas in which Clark's article (and related themes) might serve as a foundation for future directions of research. My main claims are as follows, numbered according to each section:

1. PP urges an organism-relative conception of perceptual content.
2. Historical *a priori* accounts of the structure of perceptual experience converge with results from PP.
3. There are a number of areas in which PP can find important practical applications.

Before entering into the specific issues, I should add a note about what I mean by PP. Here I am following the general theoretical framework expressed in Clark's article as well as in a number of other publications (Clark 2013; Hohwy 2013). The approach has a number of intellectual roots, including Hermann von Helmholtz (1867) and Richard Gregory (1980). The main contemporary expression of PP perhaps owes the most to Karl Friston (2005, 2008, 2010) and his collaborators, also with important developments of the generative model by Geoffrey Hinton (2007). By referring to PP as one general framework, I do not mean to imply that there are no outstanding issues of disagreement or open questions within PP. As Clark indicates, citing Spratling (2013), there are a number of options being developed as to the specific implementation of PP. Also, in the philosophical lit-

erature there is an emerging question about whether to understand PP as internalist or externalist regarding the vehicles of mental states (Hohwy 2014)—I take no position either way here, but see footnote 2. Overall, my remarks are motivated by Clark's exposition of PP, but they should be applicable to other approaches and interpretations as well.

2 A new conception of perceptual content

Clark has emphasized the way in which PP departs from the standard picture in perceptual psychology, and from David Marr's (1982) model of visual processing in particular (pp. 1–5). According to the standard account, the flow of information is “bottom-up,” as perceptual systems construct increasingly sophisticated representations based on the information transduced at the periphery. According to PP, perception involves the active prediction of the upcoming sensory input, “top-down.” Deviation from what is predicted, known as the prediction error, propagates upwards through the hierarchy until it is explained away by the Bayesian generative model.

Now I would like to add that the standard picture in perceptual psychology has been widely regarded as complementary to the standard picture in the philosophy of perception (see Tye 2000, for example). One central question in the philosophy of perception is the following: what is the *content* of perceptual states? Or, what does perception *represent*? The standard answer, in tune with Marr's approach, is that perceptual systems represent the external world, more or less as it really is. As Marr puts it, the purpose of vision is “to know what is where by looking” (1982). This way of thinking about perceptual content is almost a commonplace in the philosophical literature (Lewis 1980, p. 239; Fodor 1987, Ch. 4; Dretske 1995, Ch. 1). Kathleen Akins has described how the orthodox conception regards the senses as “servile” in that they report on the environmental stimulus “without fiction or embellishment” (1996, pp. 350–351).

Since PP overturns the reigning model in perceptual psychology, one might now ask

whether it also overturns the reigning model in the philosophy of perception. Here are two initial reasons to think that it does. First, according to PP, there is always an active contribution from the organism, or at least from a part of the organism. Perceptual states are generated internally and spontaneously by the ongoing dynamics of the generative model. Those states are *constrained* by perceptual sampling of the world, not driven by input from the world. Perceptual states are driven by the endogenous activity of the predictive brain. The relevant causal history of these states begins, if you will, within the brain, rather than from the outside. Each organism's generative model is unique in that it has been formed and continuously revised according to the particular trajectory of that organism's cycle of action and perception. As Clark himself puts it, the forward flow of sensory information is always "*relative to specific predictions*" (p. 6). These considerations make it clear that there can be variation in perceptual content for identical environmental conditions. Perceivers with different histories will have different predictions (Madary 2013, pp. 342–345). The degree of variation is an open question, but it is reasonable to expect variation.

A second reason to think that PP motivates a richer conception of perceptual content is that perception, according to PP, is not simply in the service of informing the organism "what is where." One main feature of PP is that perception and action work together in the service of minimizing prediction error. Clark explains that in "active inference [...] the agent moves its sensors in ways that amount to actively seeking or generating the sensory consequences that they [...] expect" (2013, p. 6, also see his discussion on page 16). If this is right, then perception does not serve the purpose of simply reporting on the state of the environment. Instead, perception is guided by expectation. While the received view of perceptual content answers the question of "what is out there?", PP suggests that perceptual content answers the question of "is this what I expected and tested via active inference?" In a way, PP simplifies perceptual content by replacing the goal

of representing the world with the single guiding principle of error minimization.

These two points suggest an understanding of perceptual content as something that is deeply informed by the specific history and embodiment of the organism. The content of perception is a complex interplay between particular organisms and their particular environments. At least on the face of it, this way of considering perception suggests new challenges and interesting new theoretical options for philosophers interested in describing perceptual content. For one thing, it suggests that propositional content as expressed using natural language (Searle 1983, p. 40) may be ill-suited for the task of describing perceptual content. Natural language does not typically include reports about prediction-error minimization, nor does it capture the fine-grained differences in perceptual content that will arise due to slight variations in the predictions made by different organisms. The traditional account of perceptual content, following Marr, does not include such differences, and is thus better disposed to expression using natural language.

These new challenges for understanding perceptual content may offer at the same time a general lesson for understanding all mental content in a naturalistic manner. Let me explain. One of the main goals in the philosophy of psychology has been to naturalize intentionality, to give an account of the content of mental states in terms of the natural sciences (in non-mentalistic terms). Well-known attempts include causal co-variation (Fodor 1987, Ch. 4) and teleosemantics (Millikan 1984, 2004). All attempts have met with compelling counterexamples.¹ Importantly, one implicit presupposition in the debate is that mental content should be conceived along the lines of the traditional view of perceptual content sketched above. That is, mental states are thought to be about bits of the objective world considered independently of the particular organism who possesses those mental states. To use a standard example, my belief that there is milk in the refrigerator is true if and only if there is milk in the refriger-

¹ For an overview of the major theories and their challenges, see Jacob (2010, section 9) and the references therein.

ator. This belief is about bits of the objective world: milk and the refrigerator in particular. Nothing else about my mind is deemed relevant for understanding the content of that belief. To use the familiar phrase, beliefs have a mind-to-world direction of fit (based on [Anscombe 1957](#), §32).

If my reading of PP is right, and perceptual content turns out to be a matter of the complex interaction between particular organisms and their environments, then the comfortable pre-theoretical mind/world distinction might need revision.² Recall the discussion above, in which I claimed that, on the new PP-inspired understanding of perception the question is about whether sensory stimulation fulfils the expectations of particular organisms. All perceptual states are thereby colored, as it were, by the mental lives of the organisms having those states. Organisms are not interested in what the world is like. Organisms are interested in sustaining their integrity and physical existence; they are interested in what the world is like *relative to their own particular sensorimotor trajectory through the world*, a trajectory that is partly determined by their phenotype ([Friston et al. 2006](#)). This refashioning of the mind/world relationship is unorthodox, but it is hardly new. Similar ideas can be found in [von Uexküll's Umwelt \(1934\)](#), [Merleau-Ponty's](#) discussion of sensory stimuli (1962, p. 79), [Milikan's](#) “pushmi-pullyu” representations (1995), [Akins' narcissistic sensory systems \(1996\)](#), [Clark's](#) earlier work (1997, Ch. 1), and in [Metzinger's](#) ego tunnel (2009, pp. 8–9).

Now return to the problem of naturalizing intentionality. If we replace the notion of a purely world-directed mental state with a world-relative-to-the-organism-directed mental state, then naturalizing intentionality must somehow incorporate the relationship between

the organism and its world. One way to pursue this project is to make it a matter of biology and physics. All living organisms keep themselves far from thermodynamic equilibrium by continuously exchanging matter and energy with their environment ([Haynie 2008](#)). Perhaps intentionality can be recast in terms of the organism's ongoing struggle to maintain itself as a living entity. This line of thought is central to the enactivist “sense-making” of [Maturana, Varela, and Thompson \(Maturana & Varela 1980; Thompson 2007\)](#). Crucially, it is also a central feature of [Friston's](#) version of PP. According to [Friston](#), prediction error minimization is a kind of functional description for the physical process of the organism's minimizing free energy in its effort to maintain itself far from thermodynamic equilibrium (2013). Naturalizing intentionality may be just a matter of physics (see [Dixon et al. 2014](#) for an implementation of this strategy for problem-solving tasks).

Before moving on to the next section, I should add two qualifications. First, the idea of perceptual content being partly determined by the particular history of the perceiver should not be misunderstood as some kind of radical relativism with regard to perceptual content. Even if perceptual content is *partly* determined by the details of the organism, it is also partly determined by the world itself. As proponents of PP frequently claim, our generative models mirror the causal structure of the world ([Hohwy 2013](#), Ch. 1). The point I am emphasizing here is that the causal structure of the world that is extracted is a structure relative to the embodiment (see [Clark this collection](#), section 2.4)—and perceptual history—of the perceiver. The causal structure mirrored by a chimpanzee's generative model is, in important ways, unlike the causal structure mirrored by that of a catfish.

The second qualification has to do with my remark that naturalizing intentionality may be just a matter of physics. Even if one allows that the approach I sketched shows promise, it is important to emphasize the explanatory gulf that remains. The intentionality-as-physics approach might succeed in explaining a bacterium's intentional directedness towards a

² One possibility here has been explored recently by Karl Friston using the concept of a Markov blanket, which produces a kind of partition between information states. As I read Friston, he advocates a pluralism about Markov blankets. On this view, there is not one boundary between mind and world, but instead there are a number of salient boundaries within, and perhaps around, living organisms. [Friston](#) writes that “. . . a system can have a multitude of partitions and Markov blankets . . . the Markov blanket of an animal encloses the Markov blankets of its organs, which enclose Markov blankets of cells, which enclose Markov blankets of nuclei . . .” (2013, p. 10).

sugar gradient (Thompson 2007, p. 74–75), but it is far from clear how it would apply to my belief that P—say, for example, that California Chrome won the Kentucky Derby in 2014.

The main argument of this section has been that PP motivates an understanding of perceptual content that is always organism-relative. Clark’s version of PP, while not in conflict with this idea, has not addressed it explicitly, especially as it relates to the philosophy of perception. My goal here has been to do just that.

3 Consciousness

In this section I would like to consider how conscious experience might relate to the PP framework. In particular, I suggest that there is a convergence between *a priori* descriptions of consciousness, on one hand, and the structure of information processing according to PP on the other.³ I will not remark on the way in which PP relates to some well-known issues in the study of consciousness, such as the hard problem or the explanatory gap. It is not clear to me that PP has anything new to contribute to these topics. Nor will I make any claims about which existing theories of the neural basis of consciousness fit best with PP, although I suspect there is some interesting work there to be done.

My main concern here is in the *structure* of conscious experience, of visual experience in particular. Here I adopt a strategy recommended by Thomas Nagel (1974), and David Chalmers (1996, pp. 224–225). Nagel puts the idea nicely, “[...] structural features of perception might be more accessible to objective description, even though something would be left out” (1974, p. 449, cited in Chalmers 1996, pp. 382 f.). The strategy has been implemented, in fact, using Marr’s theory of vision—the theory that, as Clark puts it, PP turns upside down. Ray Jackendoff (1987, p. 178) and Jesse Prinz (2012, p. 52) have both emphasized the structural similarities between conscious visual experience and Marr’s 2.5 dimensional sketch.

³ For a theoretical treatment of the functional significance of this convergence, see Metzinger & Gallese (2003).

Visual phenomenology is not a flat two-dimensional surface, because we see depth. But neither is visual phenomenology fully three-dimensional, because we cannot see the hidden sides of objects. Marr’s 2.5 dimensional representation captures the level in-between two and three dimensional representation that seems to correspond to our visual phenomenology; it captures Hume’s insight that visual experience is perspectival: “The table, which we see, seems to diminish, as we remove farther from it [...]” (1993, p. 104).

As Hume emphasized the perspectival nature of visual experience, Kant famously emphasized the temporal nature of experience in the second section of the *Transcendental Aesthetic*: “Time is a necessary representation (*Vorstellung*), which lays at the foundation of all intuitions” (1781/1887/1998, A31). In an elegant synthesis of these two features of visual experience, Edmund Husserl suggested that the general structure of visual experience is one of anticipation and fulfillment:

Every percept, and every perceptual context, reveals itself, on closer analysis, as made up of components which are to be understood as ranged under two standpoints of intention and (actual or possible) fulfillment. (*Logical Investigation*, VI §10 1900, Findlay trans., 1970)

In this passage from his early work, Husserl writes of “intention and fulfillment,” but he later replaced “intention” with “anticipation” when dealing with perception.⁴

The main point is fairly straightforward: we perceive properties by implicitly anticipating how the appearances of those properties will

⁴ When first developing the framework, he used the more general term “intention” because he was dealing with linguistic meaning, not perception. When applying the framework to perception one can be more precise about the nature of the empty perceptual intentions: they are anticipatory. In his later work, his *Analyses of Passive Synthesis* from the 1920s, Husserl ties in perceptual intentions with his work on time consciousness (1969) and refers to them as protentions (*Protentionen*; Husserl 1966, p. 7). In the same work, he refers to perceptual protentions as anticipations (*Erwartungen*, 1966, p. 13, and *antizipiert*, 1966, p. 7). See Madary (2012a) for a discussion of how Husserl’s framework can be situated relative to contemporary philosophy of perception. Also see Bernet et al. (1993, p. 128) and Hopp (2011).

change as we move (or as the objects move). Husserl's proposal accommodates the perspectival character of experience because it addresses the question of how we perceive objective properties despite being constrained to one perspective at a time. And it accommodates the temporal nature of experience because anticipation is always future-directed.

Here is not the place to enter into the details of the thesis that the general structure of conscious experience is one of anticipation and fulfillment (see my 2013 for some of these details), but I should add one more point. As both Husserl (1973, p. 294) and Daniel Dennett (1991, Ch. 3) have noted, peripheral vision is highly indeterminate.⁵ Also, as we explore our environment we experience a continuous trade off between determinacy and indeterminacy. As I lean in for a closer look at one object, the other objects in my visual field fade into indeterminacy. In order to account for this feature of experience, we can note that visual anticipations have various degrees of determinacy.⁶

Now let us return to PP. *If Hume provides the philosophy of perception for Marr's theory of vision, then Husserl provides the philosophy of perception for PP.* The structural similarities should be apparent. The predictive brain underlies the essentially anticipatory structure of perceptual awareness. Degrees of determinacy are encoded probabilistically in our generative models (Clark 2013; Madary 2012b). Action and perception are tightly linked (Clark this collection, p. 9) as self-generated movements stir up new perceptual anticipations.

Many readers will see a connection between the thesis of anticipation and fulfillment, on one hand, and the sensorimotor approach to perception (O'Regan & Noë 2001; Noë 2004) on the other. Overall, there is significant thematic overlap between the two (Madary 2012a, p. 149). As Seth (2014) has argued, many of the central claims of the sensorimotor approach can be incorporated into the PP framework.⁷ This synthesis offers impressive explanatory power, bringing the standard sensorimotor experimental evidence (reversing

goggles, change blindness, selective rearing) together with the theoretical neuroscience of PP. The explanatory power is even more impressive if I am correct that PP reflects the general structure of visual phenomenology, where predictive processing corresponds to perceptual anticipations and probabilistic coding corresponds to experienced indeterminacy.

4 Applied cognitive neuroscience

I would like to begin this section with some general comments about new opportunities for human self-understanding, about extending the explanandum. Academic disciplines are standardly divided into the sciences and the humanities, and some have expressed discomfort about the distance between the two modes of inquiry, or between the two cultures, as Snow (1959) famously put it (also see Brockman 1996). There is an immediate appeal to Metzinger's assertion that "Epistemic progress in the real world is something that is achieved by all disciplines together" (2003, p. 4). *If my claims from the previous section are on the right track, then we have a convergence of results between the two independent modes of inquiry, between the empirical sciences and the humanities.* It is tempting to hope that this convergence signals the beginning of a rapprochement between the sciences and the humanities. Perhaps we are at the threshold of a new science of the mind (Rowlands 2010), a science that finds natural and fruitful connections with the world of human experience. In this section, I will explore possible connections with education, public policy, and social interaction.

Clark makes two main claims in the final sections of his article that serve for the basis of my comments here. First, he suggests that PP motivates an understanding of cognitive processing as "maximally context sensitive" (p. 16), which follows from the property of PP systems being highly flexible in setting precision weightings for the incoming prediction errors. Flexibility in weighting precision enables flexibility in the deployment of processing resources. Thus there may be a wide variety of cognitive strategies at our disposal, with a continuous in-

⁵ For impressive empirical work on this theme, see Freeman & Simoncelli (2011).

terplay between more costly and less costly strategies. Second, he addresses the challenge of explaining why humans have unique cognitive powers unavailable to non-human animals who have the same fundamental PP architecture. In response to this challenge, Clark suggests that our abilities may be due to our patterns of social interaction as well as our construction of artifacts and “designer environments” (p. 19). Taken together, these two claims can be used to inform practical decisions in a number of ways.

Begin with education. Educational psychology is a broad and important area of research. PP suggests new ways of approaching human learning, ways that might depart from the received views that have guided educational psychology. I cannot begin to engage with this huge issue here, but I would like to offer one quick example. One fairly well-known application of educational psychology is in the concept of scaffolded learning, which is built on work by Lev Vygotsky and Jerome Bruner. As it is used now, scaffolded learning involves providing the student with helpful aids at particular stages of the learning process. These aids could include having a teacher present to give helpful hints, working in small groups, and various artifacts designed with the intention of anticipating stages at which the student will need help, such as visual aids, models, or tools. Clark himself mentions the abacus, which is central example of scaffolded learning (p. 19). More generally, scaffolded learning is a good example of what [Richard Menary](#) has called “cognitive practices,” which he defines as “manipulations of an external representation to complete a cognitive task” (2010, p. 238).

If PP is right, then the learning process could be optimized by designing environments in order to provide the cycle of action and perception with precisely controlled feedback (prediction error). With the growing commercial availability of immersive virtual reality equipment, educators could design learning environments (or help students design their own environments) without the messy constraints of the physical world. PP may give us a framework with which to understand—and predict—the detailed bodily movements of subjects as they

attempt to minimize their own prediction error. Using this framework, we can design systems that would optimize skill acquisition by efficiently predicting the errors that learners will make. This method could be fruitfully applied in the abstract (mathematics), the concrete (skiing), and in-between (foreign languages). Along these lines, the insights of PP, together with emerging technology, can lead to powerful new educational techniques.

Psychology is also applied in some areas of public policy. Clark mentions that PP challenges Kahneman’s well-known model of human thinking as consisting of a fast automatic system and a slower deliberative system (p. 18). Kahneman’s model has been applied as a basis for influential recommendations about laws and public policy in the United States ([Thaler & Sunstein 2008](#); [Sunstein 2014](#)). If PP homes in on a more accurate model of the thinking process, then we ought to use it, rather than (or as a complement to?) the dual systems model as a basis for policy making. Clark’s interpretation of PP suggests that we have a highly flexible range of cognitive systems, not limited to Kahneman’s two.

For example, one application of Kahneman’s model might involve the installation of environmental elements meant to appeal to the fast thinking system, to “nudge” agents towards making decisions in their best interest. If Clark is correct, we might consider even more sophisticated environmental features that have the goal of helping agents to deploy their range of cognitive strategies more efficiently. Clark’s ideas of context sensitivity and designer environments are both relevant here. As a society we may wish somehow to create environments and contexts that take advantage of the large repertoire of cognitive strategies available to us, according to Clark’s version of PP (see [Levy 2012](#), for example).

The final topic I’d like to mention in this section is what is best described in general terms as social interaction. I mean to indicate a number of related topics here, but the main issue is how PP might relate to the well-known philosophical topic of the way in which we understand and explain our behavior to one an-

other. Recall, for instance, Donald Davidson's (1963) claim that our explanation of our behavior in terms of reasons is a kind of causal explanation—reasons as causes. On his influential view, the connection between reason and actions is a causal connection. In contrast, recall Paul Churchland's envisioning of the golden age of psychology in which we dispose of folk psychological reason-giving in favor of more precise neurophysiological explanations of behavior (1981). According to Churchland's radical alternative, the causes of actions are not reasons as expressed using natural language. Instead, our actions are caused by patterns of neurons firing, patterns that can be described using mathematical tools such as a multidimensional state space. In opposition to Churchland's grand vision, we have Jerry Fodor's claim that the realization of such a vision would be "the greatest intellectual catastrophe in the history of our species" (1987, p. xii). Is PP the beginning of Churchland's grand vision coming to pass? Is a great intellectual catastrophe looming?

On one hand, PP seems like an obvious departure from folk psychology: Try explaining your X-ing to someone by claiming that you X-ed in order to minimize prediction error! One big issue here will be the way in which we think about agency itself. It seems mistaken to say that minimizing prediction error is something done by an agent. Such a process seems to be better described as occurring sub-personally. On the other hand, it is not inconceivable that propositional attitudes can capture the dynamics of prediction error minimization on a suitably coarse-grained level, perhaps along the lines suggested using symbolic dynamics (Dale & Spivey 2005; Atmanspacher & beim Graben 2007; Spivey 2007, Ch. 10). I suggest that these fascinating issues warrant further investigation. In particular, further investigation ought to incorporate Clark's ideas of maximal context sensitivity and the importance of designer environments.

The way in which we understand each other's behavior is also directly relevant for moral responsibility. Following Peter Strawson's seminal "Freedom and Resentment" (1962),

philosophers have started thinking about moral responsibility in terms of our reactions to one another, reactions that involve holding each other accountable. On one influential view, we hold each other accountable when our actions issue from our own reasons-responsive mechanisms (Fischer & Ravizza 1998). On a more recent proposal, holding each other accountable is best modeled as a kind of conversation (McKenna 2012). These proposals depend, in important ways, on assumptions about human psychology. In particular, they depend on our practice of giving reasons for behavior. As PP suggests a new fundamental underlying principle of behavior, our practices of holding each other accountable may be approached from a new perspective. The new challenge in this area will be to reconcile (if possible) the practice of giving reasons, on one hand, with PP's account of behavior in terms of error minimization on the other.

5 Conclusion

The main theme of my commentary might appear to be driven by an overexcited optimism for the new theory. To be clear, I have not claimed that PP is correct. Even its main proponents are quick to point out that important open issues remain. My claim is that it is worthwhile to consider the full implications of PP, given the convincing evidence presented so far. In this commentary, I have tried to suggest some of the implications that have not yet been mentioned—implications for perceptual content, consciousness, and applied cognitive neuroscience. These implications can be summarized as follows:

1. PP urges an organism-relative conception of perceptual content.
2. Historical *a priori* accounts of the structure of perceptual experience converge with results from PP.
3. There are a number of areas in which PP can find important practical applications.

The final section includes some challenges for future research. The main challenge is one that

has been familiar in one form or another for several decades in the philosophy of mind. This challenge is to address the tension between the way in which we understand and explain our behavior using natural language, on one hand, and our best theory of human behavior from cognitive neuroscience, which, arguably, is PP, on the other hand. In closing I should note that even if key elements of PP are eventually rejected, it might still turn out that our best model of the mind supports some of the themes I have been discussing.

Acknowledgments

I thank Thomas Metzinger and Jennifer Windt for helpful detailed comments on an earlier draft. This research was supported by the EC Project VERE, funded under the EU 7th Framework Program, Future and Emerging Technologies (Grant 257695).

References

- Akins, K. (1996). Of sensory systems and the “aboutness” of mental states. *Journal of Philosophy*, 93 (7), 337-372. [10.2307/2941125](https://doi.org/10.2307/2941125)
- Anscombe, G. E. M. (1957). *Intention*. Oxford, UK: Basil Blackwell.
- Atmanspacher, H. & beim Graben, P. (2007). Contextual emergence of mental states from neurodynamics. *Chaos and Complexity Letters*, 2 (2-3), 151-168.
- Bernet, R., Kern, I. & Marbach, E. (1993). *An introduction to Husserlian phenomenology*. Evanston, IL: Northwestern University Press.
- Brockman, J. (1996). *Third culture: Beyond the scientific revolution*. New York, NY: Touchstone.
- Chalmers, D. (1996). *The conscious mind*. Oxford, UK: Oxford University Press.
- Churchland, P. M. (1981). Eliminative materialism and the propositional attitudes. *Journal of Philosophy*, 78 (2), 67-90. [10.1111/j.1467-9973.1992.tb00550.x](https://doi.org/10.1111/j.1467-9973.1992.tb00550.x)
- Clark, A. (1997). *Being there: Putting brain, body, and world together again*. Cambridge, MA: MIT Press.
- (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Science*, 36 (3), 1-73. [10.1017/S0140525X12000477](https://doi.org/10.1017/S0140525X12000477)
- (2015). Embodied Prediction. In T. Metzinger & J. M. Windt (Eds.) *Open MIND*. Frankfurt a.M., GER: MIND Group.
- Dale, R. & Spivey, M. (2005). From apples and oranges to symbolic dynamics: A framework for conciliating notions of cognitive representation. *Journal of Experimental and Theoretical Artificial Intelligence*, 17 (4), 317-342. [10.1080/09528130500283766](https://doi.org/10.1080/09528130500283766)
- Davidson, D. (1963). Actions, reasons and causes. *Journal of Philosophy*, 60 (23), 685-700. [10.2307/2023177](https://doi.org/10.2307/2023177)
- Dennett, D. (1991). *Consciousness explained*. Boston, MA: Little, Brown, & Co.
- Dixon, J., Kelty-Stephen, D. & Anastas, J. (2014). The embodied dynamics of problem solving: New structure from multiscale interactions. In L. Shapiro (Ed.) *The Routledge handbook of embodied cognition*. London, UK: Routledge.
- Dretske, F. (1995). *Naturalizing the mind*. Cambridge, MA: MIT Press.
- Fischer, J. M. & Ravizza, M. (1998). *Responsibility and control: A theory of moral responsibility*. Cambridge, UK: Cambridge University Press.

- Fodor, J. (1987). *Psychosemantics*. Cambridge, MA: MIT Press.
- Freeman, J. & Simoncelli, E. (2011). *Metamers of the ventral stream*. . [10.1038/nrn.2889](https://doi.org/10.1038/nrn.2889)
- Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society B*, 360 (1456), 815-836. [10.1098/rstb.2005.1622](https://doi.org/10.1098/rstb.2005.1622)
- (2008). Hierarchical models in the brain. *PLoS Comput Biol*, 4 (11), e1000211. [10.1371/journal.pcbi.1000211](https://doi.org/10.1371/journal.pcbi.1000211)
- (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11 (2), 127-138. [10.1038/nrn2787](https://doi.org/10.1038/nrn2787)
- (2013). Life as we know it. *Journal of the Royal Society Interface*, 10 (86), 1-12. [10.1098/rsif.2013.0475](https://doi.org/10.1098/rsif.2013.0475)
- Friston, K., Kilner, J. & Harrison, L. (2006). A free energy principle for the brain. *Journal of Physiology - Paris*, 100, 70-87. [10.1016/j.jphysparis.2006.10.001](https://doi.org/10.1016/j.jphysparis.2006.10.001)
- Gregory, R. (1980). Perceptions as hypotheses. *Philosophical Transactions of the Royal Society B*, 290 (1038), 181-197.
- Haynie, D. (2008). *Biological thermodynamics*. Cambridge, UK: Cambridge University Press.
- Hinton, G. (2007). Learning multiple layers of representation. *Trends in Cognitive Sciences*, 11 (10), 428-34. [10.1016/j.tics.2007.09.004](https://doi.org/10.1016/j.tics.2007.09.004)
- Hohwy, J. (2013). *The predictive mind*. Oxford, UK: Oxford University Press.
- (2014). The self-evidencing brain. *Noûs*, 1-27. [10.1111/nous.12062](https://doi.org/10.1111/nous.12062)
- Hopp, W. (2011). *Perception and knowledge: A phenomenological account*. Cambridge, UK: Cambridge University Press.
- Hume, D. (1993). *An enquiry concerning human understanding*. Indianapolis, IN: Hackett Publishing.
- Husserl, E. (1900). *Logische Untersuchungen*. London, UK: Routledge.
- (1966). *Husserliana XI Analysen zur passiven Synthesis*. Dordrecht, NL: Kluwer.
- (1969). *Husserliana X zur Phänomenologie des inneren Zeitbewusstseins (1893-1917)*. Dordrecht, NL: Kluwer.
- (1973). *Husserliana XVI Ding und Raum: Vorlesungen 1907*. Dordrecht, NL: Kluwer.
- Jackendoff, R. (1987). *Consciousness and the computational mind*. Cambridge, MA: MIT Press.
- Jacob, P. (2010). Intentionality. *Stanford Encyclopedia of Philosophy, Fall*. <http://plato.stanford.edu/entries/intentionality/>
- Kant, I. (1998). *Kritik der reinen Vernunft*. Hamburg, GER: Meiner Verlag.
- Levy, N. (2012). Ecological engineering: Reshaping our environments to achieve our goals. *Philosophy and Technology*, 25 (4), 589-604.
- Lewis, D. (1980). Veridical hallucination and prosthetic vision. *Australasian Journal of Philosophy*, 58 (3), 239-249. [10.1080/00048408012341251](https://doi.org/10.1080/00048408012341251)
- Madary, M. (2012a). Husserl on Perceptual Constancy. *European Journal of Philosophy*, 20 (1), 145-165. [10.1111/j.1468-0378.2010.00405.x](https://doi.org/10.1111/j.1468-0378.2010.00405.x)
- (2012b). How would the world look if it looked as if it were encoded as an intertwined set of probability density distributions? *Frontiers in Psychology*, 3. [10.3389/fpsyg.2012.00419](https://doi.org/10.3389/fpsyg.2012.00419)
- (2013). Anticipation and variation in visual content. *Philosophical Studies*, 165, 335-347. [10.1007/s11098-012-9926-3](https://doi.org/10.1007/s11098-012-9926-3)
- Marr, D. (1982). *Vision: A computational approach*. New York, NY: Freeman and Co.
- Maturana, H. R. & Varela, F. J. (1980). *Autopoiesis and cognition: The realization of the living*. Dordrecht, NL: Reidel.
- McKenna, M. (2012). *Conversation and responsibility*. Oxford, UK: Oxford University Press.
- Menary, R. (2010). Cognitive integration and the extended mind. In R. Menary (Ed.) *The extended mind*. Cambridge, MA: MIT Press.
- Merleau-Ponty, M. (1962). *Phenomenology of perception*. London, UK: Routledge.
- Metzinger, T. (2003). *Being no one: The self-model theory of subjectivity*. Cambridge, MA: MIT Press.
- (2009). *The ego tunnel*. New York, NY: Basic Books.
- Metzinger, T. & Gallese, V. (2003). The emergence of a shared action ontology: Building blocks for a theory. *Consciousness and Cognition*, 12 (4), 549-571. [10.1016/S1053-8100\(03\)00072-2](https://doi.org/10.1016/S1053-8100(03)00072-2)
- Millikan, R. (1984). *Language, thought and other biological objects*. Cambridge, MA: MIT Press.
- (1995). Pushmi-pullyu representations. In J. Tomberlin (Ed.) *Philosophical perspectives 9: AI, connectionism, and philosophical psychology*. Atascadero, CA: Ridgeview Publishing Company.
- (2004). *Varieties of meaning: The 2002 Jean-Nicod Lectures*. Cambridge, MA: MIT Press.
- Nagel, T. (1974). What is it like to be a bat? *Philosophical Review*, 83 (4), 435-450.
- Noë, A. (2004). *Action in perception*. Cambridge, MA: MIT Press.

- O'Regan, K. & Noë, A. (2001). A sensorimotor approach to vision and visual consciousness. *Behavioral and Brain Sciences*, 24, 939-973.
- Prinz, J. (2012). *The conscious brain: How attention engenders experience*. Oxford, UK: Oxford University Press.
- Rowlands, M. (2010). *The new science of the mind*. Cambridge, MA: MIT Press.
- Searle, J. (1983). *Intentionality: An essay in the philosophy of mind*. Cambridge, MA: Cambridge University Press.
- Seth, A. (2014). A predictive processing theory of sensorimotor contingencies: Explaining the puzzle of perceptual presence and its absence in synesthesia. *Cognitive Neuroscience*, 5 (2), 97-118.
[10.1080/17588928.2013.877880](https://doi.org/10.1080/17588928.2013.877880)
- Snow, C. P. (1959). *The two cultures*. Cambridge, UK: Cambridge University Press.
- Spivey, M. (2007). *The continuity of mind*. Oxford, UK: Oxford University Press.
- Spratling, M. (2013). Distinguishing theory from implementation in predictive coding accounts of brain function. *Behavioral and Brain Sciences*, 36 (3), 51-52.
[10.1017/S0140525X12002178](https://doi.org/10.1017/S0140525X12002178)
- Strawson, P. F. (1962). Freedom and Resentment. *Proceedings of the British Academy*, 48, 1-25.
- Sunstein, C. (2014). *Why nudge?: The politics of libertarian paternalism*. New Haven, CT: Yale University Press.
- Thaler, R. & Sunstein, C. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. New Haven, CT: Yale University Press.
- Thompson, E. (2007). *Mind in life: Biology, phenomenology, and the sciences of mind*. Cambridge, MA: Harvard University Press.
- Tye, M. (2000). *Consciousness, color, and content*. Cambridge, MA: MIT Press.
- von Helmholtz, H. (1867). *Handbuch der Physiologischen Optik*. Leipzig, GER: Leopold Voss.
- von Uexküll, J. (1934). A stroll through the worlds of animals and men. In K. Lashley (Ed.) *Instinctive behavior*. New York, NY: International Universities Press.