Review of Philosophy and Psychology (2022) 13:797–808
https://doi.org/10.1007/s13164-022-00666-6

EDITORIAL

Editorial: Predictive Processing and Consciousness

Mark Miller[1] · Andy Clark[2] · Tobias Schlicht[3]

Published online: 5 December 2022

Predictive Processing (henceforth PP) is a recent, exciting framework emerging at the crossroads of cognitive science, statistical modeling and philosophy of mind (Friston 2005, 2010). Informed by recent developments in computational neuroscience and Bayesian psychology, it offers a paradigm shifting approach to studying cognition, often being presented as "the first truly unifying account of perception, cognition and action" (Clark 2015, p. 2). Its highly ambitious character is expressed in Jakob Hohwy's statement that it postulates only one mechanism which has the potential to "explain perception and action and everything mental in between" (Hohwy, 2013, p. 1). The account has already been successfully applied to a rich variety of mental phenomena, but only recently have philosophers and psychologists begun to apply it to one of the more mysterious aspects of mind, namely, consciousness. This special issue assembles some of the leading experts on the predictive processing paradigm and discusses some of its prospects and problems in this regard. In this introduction, we first sketch the explanatory framework and introduce some of the key recurring notions in this context. We then lay out some of the tasks arising from the goal of addressing consciousness with it, distinguishing those pertaining to different aspects (or kinds or concepts) of consciousness. We then provide an overview of the main ideas of the papers.

Tobias Schlicht would like to thank the Volkswagen Foundation for their financial support of his project *Situated Cognition*?(Az. 87-105) and the Thyssen Foundation for supporting a conference held at Ruhr-University Bochum in 2019 which included talks by Andy Clark, Jakob Hohwy, and Robert Rupert and where the editors had the idea to edit the present collection.

✉ Tobias Schlicht
tobias.schlicht@rub.de

1 Centre for Consciousness and Contemplative Studies, Monash University, Clayton, Australia

2 School of Media, Arts and Humanities, University of Sussex, Brighton, UK

3 Institute of Philosophy II, Ruhr-Universität Bochum, Bochum, Germany

1 Predictive Processing

Unlike more traditional approaches to neural information processing (Fodor 1975; Marr 1982), the PP framework does not assume sensory stimuli to be propagated up the cortical hierarchies before being modulated and compounded into complex representations of the world, which then guide decision making and drive behavioral responses. Rather, the framework postulates a 'top-down' approach in which the brain acts as a "sophisticated hypothesis tester" (Hohwy 2013: p. 2), constructing high-level hypotheses about the distal causes of input. These hypotheses are deployed as predictions over a wide range of cortical hierarchies, cascading down until they are tested against the patterns of peripheral activity created by the incoming stimuli. Only the differences between the predicted and the actual input (in other words, any unpredicted parts of the sensory input) get passed up the hierarchy. Depending on the estimated precision (the inverse variance) of such 'prediction error signals', the brain's model of the world is either amended and its current hypotheses are changed to accommodate the mismatch ('passive inference', perception), or the hypotheses are kept fixed and lead to resampling of the sensory states according to the current model ('active inference', action). The latter amounts to treating current hypotheses as motor commands which change the system's behavior in a way that will remove the source of the error from the environment (Hohwy 2012, p. 4).

The 'top-down' approach outlined above has been shown to be very flexible at accommodating different kinds of cognitive and perceptual phenomena (e.g., Adams et al. 2013; Feldman & Friston 2010; Rao & Ballard 1999; Spratling 2008, 2016), offering new insights into issues surrounding cognitive penetration, perceptual binding and attention (Hohwy 2013, pp. 101–138). For example, the framework has provided the most comprehensive explanation of the mechanism behind binocular rivalry to date (Hohwy et al. 2008). Binocular rivalry is a perceptual phenomenon in which subjects are simultaneously presented with two different images (e.g. a house and a face), one in each visual hemifield. Surprisingly, such a setup does not lead to an experience of an overlapping percept, but results in alternating experiences of seeing each of the images alone. The PP framework lets us derive this behavior from first principles. On this account the perceptual system is unable to resolve the rivalry between the competing hypotheses of the two images because both are equally probable given the sensory evidence – i.e., each of the hypotheses is incompatible with the sensory states in one of the visual streams, meaning that each results in equal amount of prediction error in the system. Furthermore, the system cannot synthesize these two hypotheses because they are incompatible with prior expectations about the structure of the environment (e.g., by postulating two separate objects occupying the same spatiotemporal location). Since the experimental set-up prevents any actions that could disambiguate the state of the world and the system's prior 'knowledge' prevents a formation of a mutually compatible percept, the competition is resolved in time. Although this heuristic description is quite simple, the computational model behind it (originally proposed by Hohwy et al. 2008) has been recently validated on fMRI data by Weilnhammer et al. (2017).

Despite its popularity and explanatory success, PP is facing many theoretical obstacles. For example, the question of how consciousness fits into the wider proposal of the brain being aimed at approximating hierarchically organized probabilistic inferences by minimizing prediction errors is still unanswered. After all, in our daily lives we seem to experience what J. L. Austin called "moderate-sized specimens of dry goods" and not probability density distributions (Austin 1962, p. 8). Presenting it in this way may have an intuitive appeal, but it might obscure the serious implications for the framework's viability. For example, the question of how conscious experience relates to the sub-personal processes postulated by the framework is central to elucidating in what way perceptual paradigms, such as binocular rivalry, can lend support to the framework. It also bears significance on the longstanding debates about the ontology of cognitive functions and its relationship with folk-psychological categories, as well as questions about the relation between information processing and the contents of conscious experience (Wiese 2017). Finally, remaining silent about the place of consciousness within the framework would undermine its grand explanatory ambitions, as it would mean that PP does not, after all, offer a general theory of cognition.

The success of PP in such other domains of the mental has made exploring the relationship between PP and consciousness one of the primary concerns for theorists involved in this research program. The gradual recognition of the importance of PP's relation to conscious experience has led to an increase in the number of publications on this topic, with key figures turning their attention to work on these issues. Yet, despite the awareness that the future direction of the PP research program may largely hinge on how it bridges the gap between conscious and unconscious information processing, there is little consensus about the place or role of experience in the framework. Some of the first attempts at accommodating consciousness within the framework have either been philosophically naïve (Hobson & Friston 2014, see Dolega & Dewhurst 2015 for criticism and Hobson & Friston 2016 for a reply) or confined to a very narrow subset of experiences (Seth 2014). Outlines of more general treatments exist (Hohwy 2012, Dennett 2015, Frith & Metzinger 2016), but they lack details necessary for meaningful comparison. Furthermore, such proposals are usually too general to make distinct empirical predictions which could help to test them in experimental settings.

2 Consciousness

Fortunately, a great deal of promising research on PP and consciousness has emerged in recent years. For example, Hohwy and Seth (2020) have sketched a major research program on how PP can guide the search for systematic neural correlates of consciousness in neuroscience. Clark (2019) depicts consciousness as involving multiple 'generative entanglements' in which bodily, worldly, and action-guiding information enter into complex reciprocal interactions. It is these hidden interactions, Clark argues, that deliver conscious experience by rendering the world as a meaningful arena full of opportunities and threats. Building upon this bedrock, Clark, Friston and Wilkinson (2019) argue that the broadly Bayesian perspective implied by PP can help make sense of the so-called *meta-problem of consciousness* which concerns our judgments and intuitions about consciousness as such. In particular, they suggest that

Bayesian agents with limited access to the details of their own processing regimes will inevitably infer that they possess some puzzling form of ‘qualitative awareness’, and that this undermines, or at least radically alters, the challenges involved in explaining phenomenal conscious experience (see Schlicht & Dolega 2022 for critical discussion).

Solms (2021) as well as Hohwy and Seth (2020) suggest ways in which PP and its broader cousin, the Free-Energy Principle (Friston 2005), can help dissolve the notorious *hard-problem of consciousness*, i.e., the question why any physical process should *be* or *give rise to* subjective experience at all (again, see also Schlicht & Dolega 2022). But in order to be able to systematically evaluate such proposals, we need a grasp on the character and distinctness of these *problems* and notions of consciousness. So, let’s start with some recapitulation of the problem space pertaining to the notion of consciousness. A major complication is that consciousness is not a monolithic concept but an umbrella term comprising multiple phenomena. In the more recent debate, at least the following concepts or dimensions of conscious experience have been distinguished and investigated in philosophy as well as in the empirical cognitive sciences:

a. *Phenomenal Consciousness* picks out the aspect of conscious experience that is responsible for there being something that it is like to experience (Nagel 1974). For example, if I am biting into a lemon or taking a sip of wine, these gustatory experiences involve distinctive feelings that single them out and distinguish them from others, like looking at the blue sky or suffering from a headache. Some philosophers emphasize two different aspects to phenomenal consciousness (Kriegel 2009). First, a qualitative aspect that helps individuate any given experience and separate it from others. Second, a subjective aspect that all phenomenally conscious experiences share. This aspect can be captured by pointing to the fact that what *my* conscious experiences are presenting *me* with is an object or a whole scene. That is why some philosophers take this to involve a basic sense of self-consciousness, namely, a pre-reflective self-awareness (Zahavi 2014), associated with what Ned Block (1995) calls the “me-ishness” associated with phenomenal experience. The qualitative and subjective character of phenomenal consciousness are two sides of the same coin, but they may require different explanations. The subjective characteristic of phenomenal consciousness is often described as an elusive phenomenon that has so far resisted any mechanistic explanation (Chalmers 1996) and thus provides an extraordinary challenge to any physicalist theory, in particular any reductive neuroscientific account of mental phenomena. That is why Chalmers (1996) takes it to yield a “hard problem” in contrast to several “easy” problems that have to do with cognitive functions that we also associate with consciousness. One of these functions is cognitive availability or accessibility.
b. *Cognitive Accessibility* denotes an aspect of conscious experience that has to do with the availability of sensory information for further processing and production of speech, motor control, and other forms of behavior. In contrast to the phenomenal aspect of consciousness, this is typically cast as a functional aspect of consciousness which may be easier to accommodate within a mechanistic

framework, along the lines of explanations for other cognitive functions (Block 1995, 2007). Block (2005) argues that phenomenal and access consciousness can come apart and indeed do have different and separate neural correlates, but this view is just as controversially debated as attempts to reduce consciousness, both in philosophy and cognitive science, to access consciousness. In one version, it is taken to be sufficient to explain how judgments about consciousness are generated (Dennett 1991, 2003, 2015, Cohen & Dennett 2011). In another version, illusionists about phenomenal consciousness (Frankish 2016, Kammerer 2018, Dennett 2019, Dehaene 2014) reject any notion of phenomenal consciousness that is supposed to differ from consciousness conceived of as access to information. The resulting experimental methodology relies heavily (though not exclusively) on verbal reports by participants, taken to reflect the conscious contents that have been made available to such reporting via working memory.

c. *Self and Self-consciousness.* What makes a cognitive system a self? Are there selves or is this a notion we should reject (Metzinger 2003)? Many proponents of PP rely on Metzinger's theory of a self-model to account for the phenomenal and representational aspects in informational terms within the hierarchical generative model at the core of the PP approach (Hohwy and Michael 2017, Schlicht 2017, Deane et al. 2020; see also Rupert 2022). Self-consciousness denotes the aspect of conscious experience that enables us to become aware of ourselves as subjects of experience. Notoriously, self-consciousness has also been shown to come in various degrees, ranging from a minimal pre-reflective self-awareness implicit in phenomenal consciousness (Zahavi 2014) up to very sophisticated forms of autobiographical or "narrative" self-awareness (Gallagher 2000, Dennett 1991). Many philosophers have argued for the distinction between consciousness of self as subject and consciousness of self as object, where the former plays a distinct functional role in motivating action (Shoemaker 1968, Perry 1979, Longuenesse 2017). Perry (1986) considered the sense of the self as subject the "unarticulated constituent" of an intentional state, and Hohwy and Michael (2017) already sketched a way in which the predictive processing framework could account for the self in terms of Metzinger's (2003) notion of a self-model. It is not yet settled whether this approach can only help make sense of the self as object (Schlicht & Venter 2019).

Proponents of PP must develop accounts of these various dimensions or aspects of consciousness and show how they can be related in unique ways using the framework. Now that many of the cognitive functions have been addressed in the debate and the explanatory purchase of the PP framework has been demonstrated, the time is ripe to start thinking seriously about whether and how the framework can address these central aspects of our mental life. But before we turn to the contributions to this volume, let's briefly situate PP with respect to the rival theoretical paradigms in cognitive science, namely representationalism and enactivism.

3 PP and Representationalism vs. Enactivism

PP is only the most recent game in town as far as explanatory paradigms in cognitive science are concerned. While its advocates are presenting specific accounts to a variety of mental phenomena, the representation wars are continuing in the background. The representational theory of mind (Fodor 1975), the dominant explanatory framework for decades, has been challenged by versions of enactivism (Varela et al. 1991, Thompson 2007, Hutto & Myin 2013). What is at stake is the need for mental representations in explanations of cognitive phenomena (Smortchkova et al. 2021). Enactivists argue that there is not only no need for them, they are also allegedly impossible since their content cannot be grounded naturalistically, despite optimistic efforts to provide a teleosemantic account based on Millikan's proposal (Hutto & Myin 2013, Neander 2017, Shea 2018). While Gallagher (2017) has argued that mastery of sensorimotor contingencies together with a sensitivity to affordances (action possibilities) and coupling between cognitive system and environment will replace the notion of mental representation in explanations of cognition, Clark (2015) has pointed out that some mental phenomena will remain "representation-hungry" (Clark & Toribio 1994), but that we have to rethink the familiar notion of representation and modify it in various ways. For example, in order to meet the challenges opposed by embodied cognition accounts, they must be conceived as action-oriented; and in order to fit into the PP framework, they must be conceived as probabilistic (see Lee & Orlandi 2022).

These discussions and their outcome will affect attempts of formulating a PP account of the various aspects of consciousness. In this volume, two contributions address consciousness from the enactive perspective (Gallagher, Hutto, Hipólito 2021 discuss perceptual illusions and Kiverstein, Kirchhoff and Thacker (2022) develop a PP account of pain experience).

Important questions that are addressed by the contributions to this volume include therefore, for example, whether merely the *cognitive* aspects of consciousness can be captured in this framework or whether it can also make progress in illuminating why *conscious experience* in its manifestations should be accompanied by a subjective feel and exhibit phenomenal character. Another question concerns the function of consciousness, in particular the functions of all three aspects of conscious experience, given that many proponents of the predictive processing framework argue that perception, action, attention etc. are all simply different ways of doing the same thing, namely minimizing prediction error. In what sense then could these phenomena be understood along these lines?

4 The Contributions to this Special Issue

This Special Issue features contributions by leading researchers working on the topic of PP and consciousness. By bringing together philosophical figures central to the development of the framework as well as experts on methodological approaches to the scientific study of consciousness, we hope to foster an ongoing debate which will clear off conceptual confusions and offer a unified approach to investigating consciousness within the PP framework.

Jakob Hohwy extends his account of predictive processing to the case of consciousness. Self-evidencing describes the purported predictive processing of all self-organizing systems, whether conscious or not. Self-evidencing in itself is therefore not sufficient for consciousness. Different systems may however be capable of self-evidencing in different, specific and distinct ways. Some of these ways of self-evidencing can be matched up with several properties of consciousness, and can explain them. This carves out a distinction in nature between those systems that are conscious, as described by these properties, and those that are not. This approach sheds new light on phenomenology, and suggests that some self-evidencing may be characteristic of consciousness.

Maxwell Ramstead and colleagues offer a way of applying a computational approach to the phenomenal aspect of consciousness (Ramstead et al. 2022). Their paper presents a version of neurophenomenology based on generative modelling techniques developed in computational neuroscience and biology. They call this approach *computational phenomenology* because it applies methods originally developed in computational modelling to phenomenology. Their contribution offers a new approach to neurophenomenology based on generative modelling, including an in-depth discussion of how this application of generative modelling differs from previous attempts to use it to explain consciousness. In short, generative modelling allows one to construct a computational model of the inferential or interpretive process that best explains this or that kind of lived experience.

Martina Vilas, Ryszard Auksztulewicz and **Lucia Melloni** (2022) likewise extend their computational approach to consciousness, highlighting the importance of active inference. Recently, the mechanistic framework of active inference has been put forward as a principled foundation to develop an overarching theory of consciousness which would help address conceptual disparities in the field (Wiese 2018; Hohwy and Seth 2020). For that promise to bear out, they argue that current proposals resting on the active inference scheme need refinement to become a process theory of consciousness. One way of improving a theory in mechanistic terms is to use formalisms such as computational models that implement, attune and validate the conceptual notions put forward. In this contribution, they examine how computational modelling approaches have been used to refine the theoretical proposals linking active inference and consciousness, with a focus on the extent and success to which they have been developed to accommodate different facets of consciousness and experimental paradigms, as well as how simulations and empirical data have been used to test and improve these computational models. While current attempts using this approach have shown promising results, the authors argue that they remain preliminary in nature. To refine their predictive and structural validity, testing those models against empirical data is needed i.e., new and unobserved neural data. A remaining challenge for active inference to become a theory of consciousness is to generalize the model to accommodate the broad range of consciousness explananda; and in particular, to account for the phenomenological aspects of experience. Notwithstanding these gaps, this approach has proven to be a valuable avenue for theory advancement and holds great potential for future research.

Niia Nikolova, Peter Waade, Karl Friston and **Micah Allen** (2022) turn from active to interoceptive inference within the predictive processing framework. The mainstream science of consciousness offers a few predominant views of how the brain gives rise to awareness. Chief among these are the Higher Order Thought Theory, Global Neuronal Workspace Theory, Integrated Information Theory, and hybrids thereof. In parallel, rapid development in predictive processing approaches have begun to outline concrete mechanisms by which interoceptive inference shapes selfhood, affect, and exteroceptive perception. Here, they consider these new approaches in terms of what they might offer our empirical, phenomenological, and philosophical understanding of consciousness and its neurobiological roots.

Geoffrey Lee and **Nico Orlandi** (2022) address the issue of probabilistic representations as a core element in predictive processing accounts. As mentioned above, PP construes perceptual processing as probabilistic and posits probabilistic representations. Orlandi and Lee consider three models of sensory activity from perceptual neuroscience, namely signal detection theory (SDT), probabilistic population codes (PPC), and sampling and then reflect on the sense in which the probabilistic states introduced in these models are probabilistic representations. Comparing and contrasting these probabilistic states to credences as they are understood in epistemology, they suggest that probabilistic representation, in an appropriately robust sense, can be understood as a form of analog representation. Finally, they apply this to the issue of whether conscious experience represents uncertainty, interpreting this as the claim that there are phenomenal features of experience that serve as analog probabilistic representations.

Robert Rupert (2022) turns to the topic of the self in a broadly predictive processing framework of cognitive systems and consciousness. His essay presents the conditional probability of co-contribution account of the individuation of cognitive systems (CPC) and argues that CPC provides an attractive basis for a theory of the cognitive self. The argument proceeds in a largely indirect way, by emphasizing empirical challenges faced by an approach that relies entirely on predictive processing (PP) mechanisms to ground a theory of the cognitive self. Given the challenges faced by PP-based approaches, we should prefer a theory of the cognitive self of the sort CPC offers, one that accommodates variety in the kinds of mechanism that, when integrated, constitute a cognitive system (and thus the cognitive self), to a theory according to which the cognitive self is composed of essentially one kind of thing, for instance, prediction-error minimization mechanisms. The final section focuses on one of the central functions of the cognitive self: to engage in conscious reasoning. Rupert argues that the phenomenon of conscious, deliberate reasoning poses an apparently insoluble problem for a PP-based view, one that seems to rest on a structural limitation of predictive-processing models. In a nutshell, conscious reasoning is a single-stream phenomenon, but, in order for PP to apply, two streams of activity must be involved, a prediction stream and an input stream. Thus, with regard to the question of the nature of the self, PP-based views must yield to an alternative approach, regardless of whether proponents of predictive processing, as a comprehensive theory of cognition, can handle the various empirical challenges canvassed earlier in the paper.

Two contributions address the topic from the perspective of embodied or enactive accounts of cognition. **Julian Kiverstein, Michael Kirchhoff** and **Michael Thacker** (2022) start by focusing on pain experience. Their paper aims to provide an account of the subjective character of pain experience in terms of what we will call 'embodied predictive processing'. They argue that the predictive machinery that constitutes pain experience is not brain bound but is distributed across the whole body. The prediction error minimising system that generates pain experience comprises the immune system, the endocrine system, and the autonomic system in continuous causal interaction with pathways spread across the whole neural axis. As we will see, they argue that these systems function in a coordinated and coherent manner as a single complex adaptive system to maintain homeostasis. This system, which they refer to as the neural-endocrine-immune (NEI) system, maintains homeostasis through the process of prediction error minimisation. They propose a view of the NEI ensemble as a multiscale nesting of Markov blankets that integrates the smallest scale of the cell to the largest scale of the embodied person in pain. In this way, they show how the EPP theory can make sense of how pain experience is neurobiologically constituted, and how a PP theory of pain can meet this constraint of accounting for the highly complex phenomenology of pain experience.

Shaun Gallagher, Daniel Hutto and **Inês Hipólito** (2022) argue that a number of perceptual (exteroceptive and proprioceptive) illusions present problems for predictive processing accounts. They review explanations of the Müller-Lyer Illusion (MLI), the Rubber Hand Illusion (RHI) and the Alien Hand Illusion (AHI) based on the idea of Prediction Error Minimization (PEM), and show why they fail. In spite of the relatively open communicative processes which, on many accounts, are posited between hierarchical levels of the cognitive system in order to facilitate the minimization of prediction errors, perceptual illusions seemingly allow prediction errors to rule. Even if, at the top, humans have reliable and secure knowledge that the lines in the MLI are equal, or that the rubber hand in the RHI is not our hand, the system seems unable to correct for sensory errors that form the illusion. The authors argue that the standard PEM explanation based on a short-circuiting principle doesn't work. This is the idea that where there are general statistical regularities in the environment there is a kind of short circuiting such that relevant priors are relegated to lower-level processing so that information from higher levels is not exchanged (Ogilvie & Carruthers 2016), or is not as precise as it should be (Hohwy 2013). Such solutions (without convincing explanation) violate the idea of open communication and/or they over-discount the reliable and secure knowledge that is in the system. Finally, they propose an alternative, 4E (embodied, embedded, extended, enactive) solution, arguing that PEM fails to take into account the 'structural resistance' introduced by material and cultural factors in the broader cognitive system.

Kathryn Nave, George Deane, Mark Miller and **Andy Clark** (2022) focus attention on recent developments that highlight expected future free energy.[1] They ask under what conditions the minimization of this quantity might underpin or help explain conscious experience. Their speculative suggestion is that Expected Free Energy is relevant only insofar as it delivers what Ward, Roberts & Clark (2011) have previously described as a sense of our own poise over an action space. Percep-

[1] Editor-in-chief's note: this paper, coauthored by two of the guest editors, was handled separately by executive editor F. de Vignemont as a regular submission, in agreement with the journal's guidelines. It was integrated in this SI after revisions and acceptance.

tual experience, they argue, is nothing other than the process that puts current actions in contact with goals and intentions, enabling some creatures to know the space of options that their current situation makes available. This proposal links the minimization of expected free energy to work suggesting a deep link between conscious contents and contents that are computed at an 'intermediate' level of processing, apt for controlling action.

Funding Open Access funding enabled and organized by Projekt DEAL.

References

Adams, R. A., S. Shipp, and K. J. Friston. 2013. Predictions not commands: active inference in the motor system. *Brain Struct Funct* 218 (3): 611–643.

Austin, J. L. 1962. *How to do things with words*. Clarendon Press.

Block, N. 1995. On a confusion about a function of consciousness. *Behavioral and brain sciences* 18 (2): 227–247.

Block, N. 2005. Two neural correlates of consciousness. *Trends in cognitive sciences* 9 (2): 46–52.

Block, N. J. 2007. Consciousness, Accessibility, and the Mesh between Psychology and Neuroscience. *Behavioral and Brain Sciences*, 30: 481-499. http://dx.doi.org/10.1017/S0140525X07002786.

Chalmers, D. J. 1996. *The conscious mind: In search of a fundamental theory*. Oxford University Press.

Clark, A. 2015. *Surfing uncertainty: prediction, action, and the embodied mind*. Oxford University Press.

Clark, Andy. 2019. Consciousness as generative entanglement. *Journal of Philosophy* 116 (12): 645–662.

Clark, A., K. Friston, and S. Wilkinson. 2019. Bayesing qualia: consciousness as inference, not raw datum. *Journal of Consciousness Studies* 26 (9–10): 19–33.

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

Cohen, M. A., and D. C. Dennett. 2011. Consciousness cannot be separated from function. *Trends in cognitive sciences* 15 (8): 358–364.

Deane, G., M. Miller, and S. Wilkinson. 2020. Losing ourselves: active inference, depersonalization, and meditation. *Frontiers in Psychology* 11: 539726.

Dehaene, S. 2014. *Consciousness and the brain: deciphering how the brain codes our thoughts*. Penguin.

Dennett, D. C. 1991. *Consciousness explained*. Little.

Dennett, D. C. 2003. Who's on first? Heterophenomenology explained. *Journal of Consciousness Studies, Special Issue: Trusting the Subject? (Part 1)*, 10, No. 9–10, pp. 19–30.

Dennett, D. 2015. Why and how does consciousness seem the way it seems? In: T. Metzinger, J. Windt (eds.), *Open Mind*. 10(T). Frankfurt/M.: MIND Group. doi: https://doi.org/10.15502/9783958570245.

Dennett, D. C. 2019. Welcome to strong illusionism. *Journal of Consciousness Studies* 26 (9–10): 48–58.

Dolega, K., and J. Dewhurst. 2015. Curtain call at the cartesian theater. *Journal of Consciousness Studies* 22 (9–10): 109–128.

Feldman, H., and K. J. Friston. 2010. Attention, uncertainty, and free-energy. *Frontiers in human neuroscience* 4: 215.

Fodor, J. A. 1975. *The language of thought*. MIT Press.

Frankish, K. 2016. Illusionism as a theory of consciousness. *Journal of Consciousness Studies* 23 (11–12): 11–39.

Friston, K. 2005. A theory of cortical responses. *Philosophical transactions of the Royal Society B: Biological sciences* 360 (1456): 815–836.
Friston, K. 2010. The free-energy principle: a unified brain theory? *Nature reviews neuroscience* 11 (2): 127–138.
Frith, C. D., T. Metzinger, A. K. Engel, K. Friston, and D. Kragic. 2016. What's the use of consciousness? How the stab of conscience made us really conscious. In: *The pragmatic turn: towards action-oriented views in cognitive science*. Ed. by A. Engel, K. Friston, D. Kragic. Cambridge, Mass.: MIT Press.
Gallagher, S. 2017. *Enactivist interventions: rethinking the mind*. Oxford University Press.
Gallagher, S. 2000. Philosophical conceptions of the self. *Trends in Cognitive Sciences* 4 (1): 14–21.
Gallagher, S., D. Hutto, and I. Hipólito. 2022. Predictive Processing and some disillusions about illusions. Review of Philosophy and Psychology, 1–19. https://doi.org/10.1007/s13164-021-00588-9.
Hobson, J. A., and K. J. Friston. 2014. Consciousness, dreams, and inference. *Journal of Consciousness Studies* 21 (1): 1–27.
Hobson, J. A., and K. J. Friston. 2016. A response to our Theater critics. *Journal of Consciousness Studies* 23 (3–4): 245–254.
Hohwy, J. 2012. Attention and conscious perception in the hypothesis testing brain. *Frontiers in psychology* 3: 96.
Hohwy, J. 2013. *The predictive mind*. Oxford University Press.
Hohwy, J., A. Roepstorff, and K. Friston. 2008. Predictive coding explains binocular rivalry: an epistemological review. *Cognition* 108 (3): 687–701.
Hohwy, J., and J. Michael. 2017. Why should any body have a self? In: *The subject's matter: Self-consciousness and the body*. Ed. by F. de Vignemont & A. Alsmith. MIT Press, pp. 363–391.
Hohwy, J., and A. Seth. 2020. Predictive processing as a systematic basis for identifying the neural correlates of consciousness. *Philosophy and the Mind Sciences*, *1*(II).
Hohwy, J. 2022. Conscious self-evidencing. *Review of Philosophy and Psychology*, 1-20. https://doi.org/10.1007/s13164-021-00578-x.
Hutto, D. D., and E. Myin. 2013. *Radicalizing enactivism*. MIT Press.
Kammerer, F. 2018. Can you believe it? Illusionism about consciousness and the illusion meta-problem. *Philosophical Psychology* 31: 44–67.
Kiverstein, J., M. D. Kirchhoff, and M. Thacker. 2022. An Embodied Predictive Processing Theory of Pain Experience. Review of Philosophy and Psychology, 1–26.
Kriegel, U. 2009. *Subjective consciousness: a self-representational theory*. Oxford University Press.
Lee, G., Orlandi, N. 2022. Representing Probability in Perception and Experience. *Review of Philosophy and Psychology*, 1-39. https://doi.org/10.1007/s13164-022-00647-9.
Longuenesse, B. 2017. *I, me, mine: back to Kant, and back again*. Oxford University Press.
Metzinger, T. 2003. *Being no one*. MIT Press.
Nagel, T. 1974. What is it like to be a bat? *The Philosophical Review* 83 (4): 435–450.
Nave, K., G. Deane, M. Miller, A. Clark. 2022. Expecting some action: Predictive Processing and the construction of conscious experience. *Review of Philosophy and Psychology*, 1-19. https://doi.org/10.1007/s13164-022-00644-y.
Neander, K. 2017. *A mark of the mental: in defense of informational teleosemantics*. MIT Press.
Nikolova, N., P. Thestrup Waade, K.J. Friston, M. Allen. 2022. What Might Interoceptive Inference Reveal about Consciousness? *Review of Philosophy and Psychology*, 1-28. https://doi.org/10.1007/s13164-021-00580-3.
Perry, J. 1979. The problem of the essential indexical. *Nous* 13 (1): 3–21.
Ramstead, M.J.D., J. Mago, M. Lifshitz, G. Pagnoni, R. Smith, G. Dumas, A. Lutz, K. Friston, A. Constant. 2022. From Generative Models to Generative Passages: A Computational Approach to (Neuro) Phenomenology. *Review of Philosophy and Psychology*, 1-29. https://doi.org/10.1007/s13164-021-00604-y.
Rao, R. P., and D. H. Ballard. 1999. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature neuroscience* 2 (1): 79–87.
Rupert, R. D. 2022. Cognitive Systems, Predictive Processing, and the self. *Review of Philosophy and Psychology*, 1–26. https://doi.org/10.1007/s13164-021-00577-y.
Schlicht, T. 2017. Experiencing organisms. From mineness to subject of experience. *Philosophical Studies* 175 (10): 2447–2474.
Schlicht, T., Dolega, K. 2022. You can't always get what you want. Predictive Processing and Consciousness. *Philosophy and the Mind Sciences 2*, 1-25. https://doi.org/10.33735/phimisci.2021.80.

Seth, A. K. 2014. A predictive processing theory of sensorimotor contingencies: explaining the puzzle of perceptual presence and its absence in synesthesia. *Cognitive neuroscience* 5 (2): 97–118.
Shea, N. 2018. *Representation in cognitive science*. Oxford University Press.
Shoemaker, S. S. 1968. Self-reference and self-awareness. *The journal of philosophy* 65 (19): 555–567.
Smortchkova, J., K. Dolega, T. Schlicht (eds.). 2021. *What are mental representations?* Oxford University Press.
Solms, M. 2021. *The hidden spring: a journey to the source of consciousness*. Profile books.
Spratling, M. W. 2008. Predictive coding as a model of biased competition in visual attention. *Vision research* 48 (12): 1391–1408.
Spratling, M. W. 2016. Predictive coding as a model of cognition. *Cognitive processing* 17 (3): 279–305.
Thompson, E. 2007. *Mind in life*. Harvard University Press.
Varela, F. J., E. Thompson, and E. Rosch. 1991. *The embodied mind*. Cambridge, Mass.: MIT Press.
Vilas, M., R. Auksztulewicz, L. Melloni. 2022. Active Inference as a Computational Framework for Consciousness. *Review of Philosophy and Psychology*, 1-20. https://doi.org/10.1007/s13164-021-00579-w.
Ward, D., T. Roberts, and A. Clark. 2011. Knowing what we can do: actions, intentions, and the construction of phenomenal experience. *Synthese* 181 (3): 375–394.
Weilnhammer, V., H. Stuke, G. Hesselmann, P. Sterzer, and K. Schmack. 2017. A predictive coding account of bistable perception-a model-based fMRI study. *PLoS computational biology* 13 (5): e1005536.
Wiese, W. 2017. Predictive processing and the phenomenology of time consciousness. *Philosophy and Predictive Processing.* MIND Group.
Wiese, W. 2018. Toward a mature science of consciousness. *Frontiers in Psychology* 9: 693.
Zahavi, D. 2014. *Self and other*. Oxford University Press.