

# Predictive brains and embodied, enactive cognition: an introduction to the special issue

Michael Kirchhoff<sup>1</sup>

Published online: 29 August 2017  
© Springer Science+Business Media B.V. 2017

All the papers in this special issue sit at the intersection between work on predictive processing models in the neurosciences and embodied, enactive perspectives on mind. It is arguably one of the most cutting-edge and fast-moving intersections of research in the contemporary sciences of mind and brain. All contributions deal with questions of whether and how key assumptions of the predictive brain hypothesis can be reconciled with approaches to cognition that take embodiment and enaction as playing a central and constitutive role in our cognitive lives. While there is broad consensus that bodily and worldly aspects matter to cognition, predictive processing is often understood in epistemic, inferential and representational terms. *Prima facie* this makes it hard to see how it would be possible to square embodied and enactive views, many of which are in direct opposition to inferential and representational accounts of mind, with predictive processing models. Rather than stressing how these accounts differ, others such as Clark (2016) emphasize what they have in common, focusing on how predictive processing models provide “the perfect neuro-computational partner for work on the embodied mind.” (Clark 2016, p. 1; see also Bruineberg and Rietveld 2014; Kirchhoff 2015a, b, c, 2016, 2017) In this sense, the aim of this special issue is to nudge this particular area of research forward by examining how, if possible at all, to combine the best of these frameworks in a joint pursuit of the following question: how is the mind and its enabling conditions, respectively, characterized, and how are their relations to one another best understood?

These issues were the basis of discussions that took place at a workshop that under the title ‘*Predictive Brains and Embodied, Enactive Cognition*’ was organized by Michael D. Kirchhoff at the University of Wollongong (AUS) on 6 November 2015.

---

✉ Michael Kirchhoff  
kirchhof@uow.edu.au

<sup>1</sup> Department of Philosophy, University of Wollongong, Northfields Avenue, Wollongong 2500, Australia

The discussions were both lively and productive. Most of the papers presented at this workshop are part of this special issue, with other key contributions added at a later stage.

Embodied and enactive views have dominated thinking about mind in the last decade or so, and they will most certainly continue to do so. Embodied and enactive accounts to cognition are a broad church. Yet, these approaches nevertheless gravitate around a shared set of tenets. I briefly mention some of these here. The first is a shared commitment to a rejection of traditional cognitivism and its assumption about methodological individualism. Cognitivism is the view that minds are representational in the sense that the content of such representations are understood to supervene on purely brain-based vehicles. In contrast to cognitivism, embodied and enactive cognition focus on the action-oriented character of mind, giving primacy to embodied activity over and above the re-construction of an outside world on the basis of internal, mental representations (Chemero 2009; Hutto and Myin 2013; Varela et al. 1991). An example will bring out a second tenet of the embodied, enactive take on mind. Consider the ceramist working at a potter's wheel. If one focuses on the wet clay, it becomes apparent that it slowly morphs into shape, while spinning over fast enough timescales that its shape appears to be constantly coming into being and dissipating again. The dynamic loops between the potter's body, her hands, arms, eyes, and so on, and the clay material itself seem so interwoven and fluid that the dynamic profile of the overall process is constantly forming and re-forming over short periods of time. There is, or so some argue, what Hurley terms a *dynamic singularity*: “a continuous and complex dynamic system ... with feedback loops that ... have external as well as internal orbits.” (1998, p. 333) This example, emphasizing a dynamic profile breaking across brain, body and the embedding environment, highlights the constitutive tenet of embodied, enactive views of cognition. In particular, minds are argued to be diachronically constituted in networks whose overall dynamics are distributed across neural and non-neural (bodily and worldly) parameters (Clark 2008; Engel et al. 2013; Kirchhoff 2015a, b, c, 2016; Menary 2007). The third tenet is that minds are akin to the piece of wet clay thrown onto the potter's wheel. That is, similar to the clay, minds are subject to “continuous re-shaping, re-wiring and re-modelling.” (Malafouris 2010, p. 55) In other words, minds become enculturated in virtue of the ‘plastic brain’ being shaped and transformed within a set of plastic cultural practices and scaffolded learning environments (Menary 2010a, b, 2014; Sutton 2010). Finally, this framework for thinking about cognition turns on the assumption that the specific details of human embodiment matter profoundly to the nature of mind. In short, bodies do not play a mere causal role. Instead, they play a constitutive role in the cognitive life of individual agents (Gallagher 2005; Thompson 2007).

Predictive processing is “rapidly gaining in influence” and deemed “set to dominate the science of mind and brain in the years to come” (Hohwy 2016, p. 259). Similar to the so-called E-turn in philosophy of mind and cognition—embodied, enactive and extended views—predictive processing is known under various names: the Bayesian brain hypothesis (Knill and Pouget 2004), the predictive mind (Hohwy 2013), the Cybernetic Bayesian brain hypothesis (Seth 2015), action-oriented predictive processing (Clark 2013, 2016)—all of which find a more fundamental articulation in the context of the free energy principle (Friston 2009, 2011; Friston and Stephan 2007).

In its most fundamental form, predictive processing states that organisms must act so as to maintain themselves within expected states, and that organisms can do so only by minimizing a prediction error quantity, while, at the same time, seeking to estimate and make more precise the reliability of the prediction error signal itself (Hohwy 2012). In this context, prediction error refers to the weighted mean of a random variable from probability theory and statistics. The ‘prediction error’ is a measure of a mismatch or discrepancy between the brain’s own endogenously generated predictions about the value of a specific random variable and its actual value. The term ‘precision’ is a measure of uncertainty. More precisely, it is the inverse of variance, and is in this particular sense a measure of uncertainty about a probability density. In the context of prediction error minimization, prediction errors are understood to be precision-weighted, which allows the amplitude or gain of the prediction error to depend on their expected precision (Clark 2016; Colombo and Wright 2017; Friston et al. 2012). A non-trivial assumption of predictive processing is that the brain seeks to predict the causes of its observations, and that it does so given a model of its environment, which puts the brain (or better, its predictive machinery) in a position to infer the causes of that input. This implies that minimizing prediction error is relative to a generative model. A generative model is a statistical model of how sensory input are generated. It is comprised of two components. The first is prior probability distributions over bodily and/or environmental states. And the second is generative probability distributions (or likelihoods) of an agent’s sensory input given its bodily and/or environmental causes. Utilizing its generative model, the brain is in a position to meet its impinging sensory input with predictions about the most probable causes of its sensory signals (Clark 2013; Hohwy 2012). In sum, predictive processing is the view that all psychological phenomena come about through the same process: minimization of prediction error and precision estimation. This is an explanatorily ambitious framework. For not only does it seek to explain *some* psychological phenomena under the auspices of predictive processing, it aims to explain them all. In Friston’s words, “if one looks at the brain as implementing this scheme, nearly every aspect of its anatomy and physiology starts to make sense.” (2009, p. 293) Or, as Clark says about predictive processing, it provides “a deeply unified theory of perception, cognition, and action.” (2013, p. 186) Indeed, Hohwy is equally optimistic, as he states that it explains “everything about the mind.” (2015, p. 1).

All of this makes it sound as if predictive processing is a rather unpleasant bedfellow for embodied, enactive views of mind. Talk of models, predictions, error signals, inference, etc., strikes some as too intellectualistic to coincide with talk of plasticity, embodiment, action-oriented minds, embedding environment, dynamical singularities, and coupled organism-environment systems. Yet this conclusion may be premature. Indeed, the specific aim of this special issue is to investigate whether or not these two frameworks can be seen as complementary partners in a joint pursuit of how best to understand the nature of minds and their basis.

There have been several conferences and special issues published on embodied and enactive cognitive science—e.g. Menary’s “4E Cognition: Embodied, Embedded, Enacted and Extended” (2010) in *Phenomenology and the Cognitive Sciences*. There is also increasing research on the philosophy of predictive processing, emerging with Clark’s important article “Whatever Next: Predictive Brains, Situated Agents, and

the Future of Cognitive Science” (2013) in *Behavioral and Brain Sciences*, Hohwy’s landmark monograph *The Predictive Mind* (2013) published by Oxford University Press, with Clark’s monograph *Surfing Uncertainty* (2016) also published by Oxford University Press, and Metzinger and Wiese’s Open Mind’s volume on *Philosophy and Predictive Processing*.

In the philosophical literature, we are just beginning to explore the relationship between predictive processing and embodied, enactive views. Making progress on this issue will no doubt yield substantial insights into the nature of mind. This special issue represents the first attempt at bringing the very best work on the intersection between these two paradigms together as a standalone issue in a journal volume. Thus, the present volume seeks to set the agenda for future work at this exciting crossing between predictive processing and embodied, enactive cognition.

In my opinion, there are a number of important developments in the literature on predictive processing, on the one hand, and embodied, enactive cognition, on the other, that have spurred the discussion that this special issue is concerned with and that have a substantial influence on several of the papers herein. Before giving an overview of each of the papers, I will briefly comment on a number of these developments.

The first of these developments concerns the topic of ‘*models and self-evidence*’. Over the last decade, some theorists have expressed deep scepticism about internalism with respect to the mind-world relation. Internalism can be expressed in different ways. But a common way to unpack this view is that minds are essentially internal, and “always and everywhere neurally realized.” (Clark 2008, p. xxvii) Defenders of embodied, enactive views deny internalism thus construed. Instead, they advocate for the claim that the realizers of minds and cognition are wider than the brain, comprising, at least in certain circumstances, embodiment and embedding environment. This is the constitutive thesis of embodied, enactive approaches to the mind. Nudging this way of thinking about minds as constituted in brain-body-environment networks are arguments concerning the dynamical nature of mind (Kelso 1995; Thelen and Smith 1994), sensorimotor contingency theory (Noë 2004), and the extended mind thesis (Clark and Chalmers 1998; Menary 2007; Wheeler 2010).

But there is a complication. On the assumption that the brain is involved in predictive processing, forever minimizing prediction error, prediction error minimization operates relative to a model. It suggests that minimizing prediction error is effectively to maximize model evidence, enabling an agent to successfully meet its sensory signals in a close to optimal fashion. This can also be cast as sampling the world in an epistemic sense, resolving uncertainty about future states. Yet, as we have seen, the brain does not have *direct* access to the external states causing the sensory signals it receives. On this view, for the brain to ‘know’ anything about external states, it must generate predictions about such states. Clark (2013) puts the idea as follows: “all that it [the brain] “knows”, in any direct sense, are the ways its own states (e.g. spike trains) flow and alter. In that restricted sense, all the system has direct access to is its own states. The world itself off-limits.” (2013, p. 183) Some take this to imply that minimizing prediction error is to maximize self-evidence. In other words, to increase log evidence for generative models embodied in the so-called sufficient statistics of neural architectures. What follows? Some argue that this implies that the brain is *epistemically secluded* from all other states—bodily as well as worldly states (Hohwy

2013, 2015, 2016, this issue). Others argue that it is possible to distinguish between neural generative models and embodied generative models. As Friston et al. state: “We must here understand ‘model’ in the most inclusive sense, as combining interpretive dispositions, morphology, and neural architecture, and as implying a highly tuned ‘fit’ between the active, embodied organism and the embedded environment.” (2012, p. 6; see also Bruineberg and Rietveld 2014; Bruineberg et al. 2016; Fabry, this issue; Kirchhoff 2015a, 2016, 2017) These latter ways of approaching the notion of models in predictive processing push for a notion of models that is more attuned with work in embodied, enactive cognitive science. Of course, this discussion is far from settled. The jury is still out. In fact, a third alternative is to start from the free energy principle, and from there attempt to illustrate how the free energy principle itself alleviates any tension between internalist and externalist accounts of mind (Allen and Friston, this issue).

The second of these developments concerns the vexed problem of ‘*internal mental representation*’. On some radical formulations of embodied, enactive cognition, there is no justified need for the positing of internal mental representations to explain cognition (or minimally, most forms of cognitive activity). It is a claim that runs counter to mainstream cognitive science. One might try to underpin this non-representationalist claim along the following lines. Suppose that cognitive systems are reciprocally coupled brain-body-environment systems. If so, then there is no need for one part (the organism) to represent another part (the world). As Silberstein and Chemero put it: “if the animal-environment system is just one system, the animal portion of the system need not represent the environment portion of the system to maintain its connection with it.” (2012, p. 40) Silberstein and Chemero (2012) are not alone. Further bolstering this view is the idea that minds are *not* primarily for thinking, traditionally conceived, but *for action* (Hutto and Myin 2013). Indeed, as Froese and Ziemke mention, these radical E-approaches have “matured and become a viable alternative” to representationalist cognitive science resulting in “methodological advances” that “avoid or successfully address many of the fundamental problems” faced by representationalist accounts of mind (2009, p. 466).

Yet, predictive processing appears committed to representationalism. Consider, e.g., what Clark says here:

There remains, however, at least one famously vexed issue upon which PP and the enactivist ... seem doomed to disagree. That is the issue of ‘internal representation.’ Thus Varela et al. are explicit that, on the enactivist conception “cognition is no longer seen as problem solving on the basis of representations” [1991, p. 205]. PP, however, deals extensively in internal models – rich, frugal, and all points in-between – whose role is to control action by predicting complex plays of sensory data. This, the enactivist might fear, is where our promising story about neural processing breaks bad. (2015, pp. 20–21)

This problem divides the waters. Some argue that predictive processing can be seen through the lens of embodied, enactive cognition, and that this excludes the need for assuming that predictive processing must be representational (Bruineberg et al. 2016; Hutto, this issue; Kirchhoff 2015a; Orlandi 2014, this issue). Others endorse

the representationalist thesis of predictive processing on the basis of adopting the Helmholtzian account of perception as unconscious inference (Hohwy 2013), while others again focus on revamping the notion of representation—taking representation to be probabilistic, action-oriented and as not concerned with conditions of truth and falsity (Clark 2015; Orlandi 2014; see also Gładziejewski 2015). We will see this discussion unfolding in some of the articles comprising this special issue.

The third development I wish to touch upon concerns the ‘*development of phenotypic traits*’. The movement of embodied, enactive cognition has a rich and varied set of tools at its disposal when it comes to addressing phenotypic development across both phylogenetic and ontogenetic timescales. One example is the concept of niche construction, a phenomenon often referred to in this literature. According to its developers, niche construction is “the process whereby organisms modify their own and/or each others’ niche, through their metabolism, their activities, and their choices.” (Laland et al. 2009, p. 196) Examples of niche construction include nest manufacturing in, e.g., birds, burrows, holes and webs, as well as houses, clothing, technologies, and so on. Niche construction is concerned with the “causal basis of the organism-environment match and its inherent symmetries, and the active role that organisms play in driving evolutionary and co-evolutionary events.” (Laland et al. 2009, p. 196).

Predictive processing, by contrast, it is just more of the same: minimizing prediction error. Indeed, for predictive processing, the whole function of the brain, including an organism’s phenotypic traits, is prediction error minimization. This singular emphasis has led some to be sceptical about the explanatory power of predictive processing when it comes to addressing the development of phenotypic traits. This species of scepticism is often cast as the so-called Dark Room problem (Clark 2016; Friston et al. 2012; Hohwy 2013). It goes like this. If the sole function of an organism is to minimize prediction error, then organisms should—in virtue of this imperative—occupy states with the lowest possible error signal. And where else could that be than in a wholly predictable environment? In other words, where else could that be than in the least surprising and, by extension, least stimulating environment? If correct, the Dark Room problem highlights that predictive processing per se cannot explain what fuels and drives adaptive success. For its fails to account for what it is that enables phenotypic development (Menary 2015; Klein, this issue). Some concede this point, yet think that predictive processing can meet this problem if appropriately complemented by cultural practices, nested and dynamical organism-niche interactions, and so on (Fabry, this issue). Others, like Kirchhoff (2016) focuses on the free energy principle, from which he argues that the Dark Room problem can be shown to be not plausible as an objection to the free energy principle.

These three developments signify that the purpose of this special issue to get clearer about how best to understand the relation between predictive processing models and embodied, enactive cognition.

Yet this list of developments is not exhaustive. To see this, consider that the current literature on predictive processing has focused almost exclusively on *cortical* aspects of hierarchical predictive processing, on the one hand, and has yet to fully address how emotions fit into this emerging picture of the brain as a predictive engine, on the other (for initial discussion, see Hohwy 2011, 2013; Seth and Friston 2016). This raises a couple of important questions. The first is: how, if at all, does emotions and sub-cortical

processes play a role in predictive processing? The second question is: if both emotions and sub-cortical processes have a role to play, what is the nature of such contributions? Is it merely a causal role, or is it a constitutive role. Both of these questions are addressed by Miller and Clark (this issue). Other outstanding questions concern the relationship between embodiment, dreaming and predictive processing (Windt, this issue), and how (if at all) the predictive brain hypothesis informs discussions of social cognition in the context of active inference (Gallagher and Allen, this issue).

Let us thus see next what the contents of the papers are:

The first paper in this special issue ‘Predictive perceptual systems’ is by **Nico Orlandi**. She starts by arguing that theories of predictive processing are sometimes treated as a monolithic block, taken to support perceptual Bayesian inference and certain intellectualist and Kantian ways of looking at perception. She ascribes this view to **Hohwy (2013)**. In this paper, Orlandi offers a systematic treatment of these assumptions, eventually problematizing this ‘attitude’ by distinguishing between different strands of predictive processing. In particular, she looks at the complex relationship between Bayesian accounts, on the one hand, and both Empirical Ranking Theory and Natural Scene Statistics, on the other. The negative result of this treatment is scepticism about (a) the pairing of predictive processing with Bayesian inference, (b) the intellectualist interpretation of predictive processing, and (c) the idea that predictive systems are veridical systems. The positive result is that from the point of view of predictive processing, perceptual systems can be understood as embedded and ecological systems skewed by organismic needs.

**Alex Kiefer** and **Jakob Hohwy** turn to consider how content arises within the prediction error minimization framework. They start by presenting the problem of content, and then propose how the generation of unsupervised representations is possible in prediction error minimization architectures, with access only to its internal and sensory states. They develop this case by exploring the possibility of misrepresentation, and they explain how misrepresentation is possible in terms of the formal and probabilistic tools put to use in the prediction error minimization scheme (e.g., the Kullback-Leibler divergence between the exact and the approximate probability densities). In the final parts of their contribution, Kiefer and Hohwy pause to reflect on potential limitations to their proposal, while discussing specific aspects of representational systems in the light of prediction error minimization.

The dispute about the representational status of predictive processing continues in the contribution by **Jelle Bruineberg**, **Julian Kiverstein** and **Erik Rietveld** entitled ‘The brain is not a scientist: the free-energy principle from an ecological-enactive perspective’. Unlike the first two articles, Bruineberg, Kiverstein and Rietveld push the separation of the free energy principle from Helmholtzian accounts of the predictive brain. They start by observing that the free energy principle has typically been connected with a Bayesian theory of predictive coding, and that the latter is often taken to support a Helmholtzian theory of perception as unconscious inference. However, as they argue, approaching Bayesian predictive coding from the perspective of the free energy principle is not compatible with a Helmholtzian account of perception. This ultimately leads the authors to defend and argue for three claims. The first is that free energy minimization is widely realized. Thus, the processes of free energy minimization are argued to involve the entire animal-environment system. The second is that

the notion of active inference, as it is used in the free energy principle, is not easy to shoehorn with unconscious inference understood through hypothesis testing. The third and final argument targets the notion of inference. They argue that the notion of inference is much too weak in Bayesian predictive processing to support anything like the view that it turns on hypothesis testing. Taken together, these arguments are said to imply that the free energy principle is best understood in ecological and enactive terms.

In his ‘Getting into the great guessing game: Bootstrap heaven or hell?’ **Daniel D. Hutto** defends a radically enactive take on predictive processing, thus bringing his trademark radical enactivism to bear on the fortunes of predictive processing. Hutto’s focus can be cast in the form of a continuity argument. The first thing that he points to is that predictive processing, in what he calls its cognitivist rendering, understands perceiving in terms of guessing. To perceive the world is to guess the world, as he says. The question that Hutto raises is how the brain is able to get into the guessing game in the first place. A different way of articulating this question would be to ask: where do the priors come from? In a close treatment of Clark’s (2016) solution to this issue, one Clark dubs ‘bootstrap heaven’, Hutto argues that once a realistic look at the situation of the brain is provided, Clark’s bootstrap heaven collapses to a bootstrap hell. The positive contribution of this paper is that Hutto goes on to suggest a solution to this problem, one that implies adopting a radical enactive take on predictive processing.

The final contribution of this first cluster of papers focusing (broadly speaking) on the implications of predictive processing for boundaries of mind discussion is by **Micah Allen and Karl Friston**. In their, ‘From cognitivism to autopoiesis: Towards a computational framework for the embodied mind,’ they start by giving a review of neuroscientific, cognitive and philosophical approaches to predictive processing. They do so in order to illustrate how these different approaches sit on a continuum ranging from cognitivist commitments to modular, internalistic views embracing the need for representational explanation, to more moderate embodied views emphasizing the significance of so-called ‘body-representations’ to accounts that connects with the recent radical enactive, embodied and dynamic movement in theories of mind. From here they go on to argue for two things. The first is that any nascent theory of predictive processing must consider this continuum of views, including their implicit or explicit theoretical commitments. The second is that by starting from the free energy principle, it is possible to alleviate the tension between internalist and externalist views of cognition by showing how a formal synthetic account of internal representations arise from processes of autopoietic self-organization.

This brings us to the next cluster of papers. In ‘Betwixt and between: The enculturated predictive processing approach to cognition’ **Regina E. Fabry** argues that by itself the predictive processing framework is not as explanatorily powerful as it could be. That is, as a framework that applies purely to the brain, it cannot provide a coherent and plausible account of a broad class of complex, socio-culturally structured cognitive processes. Following in the steps of **Menary (2007)**, she coins these processes ‘cognitive practices’. Fabry argues that if predictive processing accounts are to have such explanatory power, they must (necessarily) be complemented by accounts of the nested and dynamical nature of organism-niche interactions given that only such a complementary account can capture and explain the nontrivial role of



enculturation, phylogenetic as well as ontogenetic influences on cognitive practices. Consequently she pursues and defends what she coins the *enculturated predictive processing* framework—one capable of unifying predictive profiles of organisms across individual and environment.

The contribution by **Michael D. Kirchhoff**, ‘Autopoiesis, free energy and the life-mind continuity thesis,’ targets a strong form of the life-mind continuity thesis (see also [Kirchhoff and Froese 2017](#)). It does so by providing a critical comparison of the theory of autopoiesis (developed in biology and enactivist theories of mind) and the free energy principle (with roots in thermodynamics and statistical physics). Kirchhoff advances two claims. The first is that the free energy principle should be preferred to theory of autopoiesis, as classically formulated, when targeting the life-mind continuity thesis. One reason for this is that the free energy principle is based on an externalist version of the causal-explanatory relation between organism and environment, whereas the classical theory of autopoiesis turns on an internalist kind of this explanatory relation. From this observation, Kirchhoff goes on to argue that only an externalist kind of this relation is able to accommodate adaptivity, which is shown to be at the core of both life and mind. Another reason is that the theory of autopoiesis, as classically formulated, is open to the Dark Room problem, while Kirchhoff argues that the free energy principle is immune to this problem. The second claim is that the free energy principle and the recently formulated framework of autopoietic enactivism can be shown to be continuous on a number of central issues, thus raising the possibility of a joint venture when it comes to answering the life–mind continuity thesis.

**Colin Klein’s** contribution ‘What do predictive coders want?’ ends this second group of papers. He also focuses on the Dark Room problem. Yet he does so from a different starting point. Klein argues the Dark Room problem highlights the challenges that predictive processing theories face in accounting for motivational states. If Klein is right, then none of the extant responses work, and they fail on principled grounds. He further argues that the free energy principle and the Good Regulator Theorem are covering laws. According to Klein such laws cannot explain motivational states, and therefore cannot solve the Dark Room problem. As he points out, real-world and real-time systems embody motivational states in various ways consistent with these large-scale covering laws. Indeed, they do so in ways that are non-representational. This leads Klein to the conclusion that predictive processing can reconcile its account of mind with more embodied principles, even if it ultimately undermines their deeper explanatory ambitions.

**Mark Miller and Andy Clark’s** contribution opens the third and final group of papers. In their ‘Happily entangled: Prediction, emotion, and the embodied mind,’ they start by making the observation that contemporary work in cognitive and computational neuroscience depict that human cortex as a multi-level prediction engine. Like others in the field, they also think that the predictive processing framework shows great promise when it comes to understanding and integrating essential information processing strategies underlying perception, reasoning, and action. On the basis of this they ask: how, if at all, emotions and sub-cortical contributions fit into this emerging view of the brain as a predictive organ? Miller and Clark provide a positive answer.

They argue that the ‘fit’ is profound and potentially transformative. Emotions and sub-cortical processing do not operate separate from the multi-level predictive core. Instead, it is argued, emotion and sub-cortical processing appear to be inextricably bound up with the construction of multi-level predictions that determine what we perceive and how we act. For Miller and Clark, these entanglements suggest new and more satisfying responses to several common worries concerning the explanatory power of the predictive processing framework.

**Jennifer Windt** focuses on dreaming from which she draws two conclusions. She starts by considering that if dreaming depends on brain activity alone, and if consciousness is modelling on dreaming, then conscious experience is realized by brain activity alone. This is internalism. Yet one need not accept this implication. As Windt mentions, proponents of embodied, enactive cognition often accept that dreaming is a state of functional disembodiment, while, at the same time, denying that the same is the case for conscious experience. In her contribution, Windt seeks to dampen the appeal of the view that dreaming is an example of what she calls a naturally occurring cranial envatment. She does this by arguing that dreams are weakly phenomenally-functionally embodied states in which the processing of peripheral and bodily stimuli is altered, without being entirely suppressed. According to Windt, the result of this is that the cranial envatment view of dreams fails, thus depriving internalist of their most significant empirical example. From here Windt turns to consider how predictive processing accounts of dreaming can be enriched through a focus on the phenomenal-functional configuration that underlies self-experience and self-other distinctions in dreams.

In ‘Active inference, enactivism and the hermeneutics of social cognition,’ **Shaun Gallagher** and **Micah Allen** target social cognition, the final frontier to be tackled in this special issue. They start by distinguishing between three different philosophical views on the neuroscience of predictive models: predictive coding (associated with Jakob Hohwy (2016, this issue), predictive processing (associated with Andy Clark (Clark 2015, 2016; this issue) and what they call predictive engagement (associated with enactivist approaches to cognition, see e.g., Bruineberg and Rietveld 2014; Bruineberg et al. 2016; Kirchoff 2015a, b, this issue). For each of these different philosophical predictive models, Gallagher and Allen investigate the concept of active inference and ask how, if at all, this concept can inform work on and our best understanding of social cognition. Building on Friston and Frith’s (2015 notion of neural hermeneutics, Gallagher and Allen offer an alternative model of enactivist hermeneutics, which they argue supports a predictive engagement account of social cognition.

**Acknowledgements** Kirchoff’s work was supported by an Australian Research Council Discovery Project “Minds in Skilled Performance” (DP170102987), a John Templeton Foundation grant “Probabilizing Consciousness: Implications and New Directions” hosted at the University of Cambridge as part of the New Directions in the Study of the Mind project, and by a John Templeton Foundation Academic Cross-Training Fellowship (ID#60708). The opinions expressed in this publication are those of the author and do not necessarily reflect the views of the John Templeton Foundation. Thanks also to an anonymous reviewer for helpful comments. Special thanks to all the reviewers involved in the entire special issue without whom this project would not have seen its realization. Thanks to the editors at Synthese, especially to Wiebe Van Der Hoek for editorial assistance along the way. And finally a huge thanks to all the contributing authors.

## References

- Bruineberg, J., Kiverstein, J., & Rietveld, E. (2016). The brain is not a scientist: The free-energy principle from an ecological-enactive perspective. *Synthese* 1–18.
- Bruineberg, J., & Rietveld, E. (2014). Self-organization, free energy minimization, and optimal grip on a field of affordances. *Frontiers in Human Neurosciences*, 1–14. doi:10.3389/fnhum.2014.00599.
- Chemero, A. (2009). *Radical embodied cognitive science*. Cambridge, MA: The MIT Press.
- Clark, A. (2008). *Supersizing the mind: Embodiment, action and cognitive extension*. Oxford: Oxford University Press.
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36, 181–253.
- Clark, A. (2015). Radical predictive processing. *The Southern Journal of Philosophy*, 53(S1), 1–25.
- Clark, A. (2016). *Surfing uncertainty*. Oxford: Oxford University Press.
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58, 10–23.
- Colombo, M., & Wright, C. (2017). Explanatory pluralism: An unrewarding prediction error for free energy theorists. *Brain and Cognition*, 1–10. doi:10.1016/j.bandc.2016.02.003.
- Engel, A. K., Maye, A., Kurthen, M., & König, P. (2013). Where's the action? The pragmatic turn in cognitive science. *Trends in Cognitive Science*, 17(5), 202–209.
- Friston, K. (2009). The free-energy principle: A rough guide to the brain? *Trends in Cognitive Sciences*, 13(7), 293–301.
- Friston, K. (2011). Embodied inference: Or 'I Think Therefore I Am, If I Am What I Think'. In W. Tschacher & C. Bergomi (Eds.), *The implications of embodiment (cognition and communication)* (pp. 89–125). Exeter: Imprint Academic.
- Friston, K., & Frith, C. (2015). Active inference, communication and hermeneutics. *Cortex*, 68, 129–143.
- Friston, K., & Stephan, K. E. (2007). Free energy and the brain. *Synthese*, 159, 417–458.
- Friston, K., Thornton, C., & Clark, A. (2012). Free-energy minimization and the dark-room problem. *Frontiers in Psychology*, 3(130), 1–7.
- Froese, T., & Ziemke, T. (2009). Enactive artificial intelligence: Investigating the systemic organization of life and mind. *Journal of Artificial Intelligence*, 173, 466–500.
- Gallagher, S. (2005). *How the body shapes the mind*. Oxford: Oxford University Press.
- Gładziejewski, P. (2015). Predictive coding and representationalism. *Synthese*, 193(2), 559–582.
- Hobson, A., & Friston, K. (2014). Consciousness, dreams, and inference. *Journal of Consciousness Studies*, 21(1–2), 6–32.
- Hohwy, J. (2011). Phenomenal variability and introspective reliability. *Mind & Language*, 26(3), 261–286.
- Hohwy, J. (2012). Attention and conscious perception in the hypothesis testing brain. *Frontiers in Psychology*, 1–4. doi:10.3389/fpsyg.2012.00096.
- Hohwy, J. (2013). *The predictive mind*. Oxford: Oxford University Press.
- Hohwy, J. (2015). The neural organ explains the mind. In T. Metzinger & J. M. Windt (Eds.), *Open mind* (Vol. 19, Issue 1, pp. 1–22). Frankfurt am Main: MIND Group.
- Hohwy, J. (2016). The self-evidencing brain. *Nous*, 50(2), 259–285.
- Hurley, S. L. (1998). *Consciousness in action*. Cambridge, MA: Harvard University Press.
- Hutto, D. D., & Myin, E. (2013). *Radicalizing enactivism: Basic minds without content*. Cambridge, MA: The MIT Press.
- Kelso, S. (1995). *Dynamic patterns*. Cambridge, MA: The MIT Press.
- Kirchhoff, M. D. (2015a). Experiential fantasies, prediction, and enactive minds. *Journal of Consciousness Studies*, 22(3–4), 68–92.
- Kirchhoff, M. D. (2015b). Species of realization and the free energy principle. *Australasian Journal of Philosophy*, 93(4), 706–723.
- Kirchhoff, M. D. (2015c). Extended cognition & the causal-constitutive fallacy: In search for a diachronic and dynamical conception of constitution. *Philosophy and Phenomenological Research*, 90(2), 320–360.
- Kirchhoff, M. D. (2016). Autopoiesis, free energy, and the life-mind continuity thesis. *Synthese*. doi:10.1007/s11229-016-1100-6.
- Kirchhoff, M. D. (2017). Predictive processing, perceiving and imagining: Is to perceive to imagine, or something close to it? *Philosophical Studies*. doi:10.1007/s11098-017-0891-8.
- Kirchhoff, M. D., & Froese, T. (2017). Where there is life there is mind: In support of a strong life-mind continuity thesis. *Entropy*, 19(4), 169. doi:10.3390/e19040169.

- Knill, D. C., & Pouget, A. (2004). The Bayesian brain: The role of uncertainty in neural coding and computation. *Trends in Neuroscience*, 27(12), 712–719.
- Laland, K. N., Odling-Smee, J., Feldman, M. W., & Kendal, J. (2009). Conceptual barriers to progress within evolutionary biology. *Foundations of science*, 14, 195–216.
- Malafouris, L. (2010). Metaplasticity and the human becoming: Principles of neuroarchaeology. *Journal of Anthropological Sciences*, 88, 49–72.
- Menary, R. (2007). *Cognitive integration: Mind and cognition unbounded*. Basingstoke: Palgrave.
- Menary, R. (2010). Introduction to the special issue on 4E cognition. *Phenomenology and the Cognitive Sciences*, 9, 459–463.
- Menary, R. (2014). Neural plasticity, neuronal recycling and niche construction. *Mind & Language*, 29(3), 286–303.
- Menary, R. (2015). Predictive coding and enculturation: A reply to Fabry. In T. Metzinger, J. Windt (Eds.), *Open MIND: 25(T)*. Frankfurt am Main: MIND Group. doi:10.15502/9783958570818.
- Noë, A. (2004). *Action in perception*. Cambridge, MA: The MIT Press.
- Orlandi, N. (2014). *The innocent eye: Why vision is not a cognitive process*. Oxford, NY: Oxford University Press.
- Seth, A. K. (2015). The cybernetic Bayesian brain. In T. Metzinger & J. M. Windt (Eds.), *Open MIND* (Vol. 35, Issue 1, pp. 1–24). Frankfurt am Main: MIND Group.
- Seth, A. K., & Friston, K. (2016). Active interoceptive inference and the emotional brain. *Philosophical Transactions B: Biological Sciences*, 371, 20160007.
- Silberstein, M., & Chemero, A. (2012). Complexity and extended phenomenological-cognitive systems. *Topics in Cognitive Science*, 4, 35–50.
- Thelen, E., & Smith, L. (1994). *A dynamic systems approach to the development of cognition and action*. Cambridge, MA: The MIT Press.
- Thompson, E. (2007). *Mind in life: Biology, phenomenology, and the sciences of mind*. Cambridge, MA: Harvard University Press.
- Varela, F., Thompson, E., & Rosch, E. (1991). *The embodied mind*. Cambridge, MA: The MIT Press.
- Wheeler, M. (2010). In defense of extended functionalism. In R. Menary (Ed.), *The extended mind*. Cambridge, MA: The MIT Press.