What? Now. Predictive Coding and Enculturation

A Reply to Regina E. Fabry

Richard Menary

Regina Fabry has proposed an intriguing marriage of enculturated cognition and predictive processing. I raise some questions for whether this marriage will work and warn against expecting too much from the predictive processing framework. Furthermore I argue that the predictive processes at a sub-personal level cannot be driving the innovations at a social level that lead to enculturated cognitive systems, like those explored in my target paper.

Keywords
Active inference | Cognitive integration | Enculturation | Learning driven plasticity | Mathematical cognition | Perceptual inference | Predictive processing | Reading

Author

Richard Menary
richard.menary@mq.edu.au
Macquarie University
Sydney, NSW, Australia

Commentator

Regina E. Fabry
fabry@students.uni-mainz.de
Johannes Gutenberg-Universität
Mainz, Germany

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: What? Now.

I'd like to thank Regina Fabry for her excellent and detailed response to my paper. She articulates an important account of reading acquisition as a process of enculturation and describes how a Cognitive Integration/Enculturated Cognition (henceforth CI/ENC) account can be combined with a predictive processing account of neural processing. She shows, in impressive detail, how CI/ENC can benefit from Predictive Processing (henceforth PP), primarily as a way of explaining the neural-level details of processes that conspire with bodily interactions with the local environment to complete cognitive tasks. Since Fabry's response suggests an important way of cashing out some of the details of an enculturated approach, I would like to take this opportunity to look at some of the potential pitfalls in the proposed Enculturated Predictive Processing style (henceforth EPP). Primarily I want to focus on the differences in explanatory emphasis between CI/ENC and PP, especially where CI/ENC proposes the importance of the population-level effects of normative patterned practices (henceforth NPP), such as mathematical practices.

PP is all about predictions, happening in the here-and-now[1]; however CI/ENC occurs at different levels and over much longer time-scales. It turns out that this difference is important, because if the brain is engaged in predictive error minimization (as sub-personal processing) in the here-and-now, then it cannot be driving the innovation of new NPP over many generations. This is because the pressures driving those innovations are found at the social, or populational, level[2], not at the level of neural processing where 'what?' is answered in the now.

I also raise several issues concerning the nature of the PP project, particularly whether, as a theory of general brain architecture, all processing can be cashed out in terms of predictive processes. I'm also sceptical about Fabry's claim that PP can provide the "mechanistic underpinnings of the acquisition of cognitive practices" (Fabry this volume, p. 3) on its own, without help from what I call learning-driven plasticity (LDP) and neural redeployment. Finally, I comment on the promising research path down which Fabry is headed.

In the first section I remind the reader of some of the leading ideas of the CI/ENC framework, highlighting, in particular, the different levels of explanation and how this matters to the proposed marriage of ENC-PP. In the second section I raise several problems for the PP approach in general and for the ENC-PP approach in particular. My main concerns are to push away from an 'isolated brain' interpretation of PP and to place EPP within a much broader context of explanation.

2 CI and enculturation

As I point out in my contribution to this volume, cognitive integration should be understood as a thesis about the enculturation of human cognition. It is a thesis about how phylogenetically earlier forms of cognition are built upon by more recent cultural innovations (e. g., systems of symbolic representation). This results in a multi-layered system with heterogeneous components, dynamically interwoven into a co-operative of processes and states an integrated cognitive system (henceforth ICS). The co-ordination dynamics of the system are, at least in part, understood in terms of the physical dynamics of brain–body–niche interactions in real-time; however, they are also to be understood in terms of NPP that govern and determine those interactions (over time). NPP operate at both social/population levels and individual, even sub-personal, levels. They originate as patterns of activity spread out over a population of agents; consequently they should be understood primarily as public systems of activity and/or representation that are susceptible to innovative alteration, expansion, and even contraction over time. They are transmitted horizontally across generational groups and vertically from one generation to the [3]next. At the individual level they are acquired, most often by learning and training, and they manifest themselves as changes in the ways in which individuals think, but also the ways that they act (intentionally) and the ways in which they interact with other members of their social group(s) and the local environment. NPP, therefore, operate at different levels (groups and individuals) and over different time-scales (intergenerationally and in the here-and-now).

Given this, it is clear that What? Now[4] processes that reduce prediction errors on their own could not drive the innovation of NPP; nor could they determine the properties of NPP on their own. Less obviously, I would argue, they do not drive the acquisition of NPP, because scaffolded learning requires both a physically and temporally-structured learning environment and the capacity for functional changes to cortical circuitry to be driven by the structured learning environment. The mechanism of acquisition includes both neural and environmental processes working in concert and over long periods of ontogenetic time. What? Now processes may help us to understand the here-and-now

1 I mean predictions on incoming sensory input relevant to immediate action in the environment.

2 I think that these levels are real. There is a level of entire populations, social groups, individual organisms and there is a level of individual brains. Cognition takes place within and across (at least the final three) levels.

3 See my target article for examples.

4 Predictions on sensory input in the here-and-now.

processes by which we enact NPP; they may even tell us something about the neural mechanisms for learning and plasticity; but we should be wary of making prediction and error minimization the driving factors behind the why and how of enculturation.

Fabry's commentary focuses on the neural level, functioning in real-time, where the primary aim is to give a mechanistic account of how cognitive capacities can be transformed by learning and training in rich socio-cultural niches. Rather than looking at the origin of ICS in cultural inheritance, phenotypic plasticity, and learning driven plasticity, Fabry argues that a version of the PP framework can provide the neural mechanisms by which ICS are (partly) constructed. My contribution to this volume focused primarily on the origin of ICS in the recent cultural evolution of NPP and then explored how mathematical practices could be learnt and how this process of learning could drive functional changes to circuitry in the brain. Consequently, the CI/ENC framework pursues the phylogenetic and ontogenetic basis of the larger brain–body–niche nexus. What, though, of the neural mechanisms of transformation?

I don't agree with Fabry's starting premise that CI/ENC lacks a mechanism of transformation: the mechanism of transformation is learning-driven plasticity (LDP) with neural redeployment in a scaffolded learning environment. The fundamental plasticity of the brain explains the nature of neural transformations and why the brain is open to scaffolded learning driven by the environment. (E)PP does not have the resources to explain redeployment (this is a theme I take up in the next section). Why would it, since PP is not a framework for explaining redeployment. It might be the case that PP fits with a certain conception of scaffolded learning such as path-dependent learning, but I have yet to see a thorough working-through of the details and it's not clear to me that all scaffolded learning should be reduced to a predictive form of path-dependent learning.

Fabry claims that a dynamical systems approach to integration "does not spell out the mutual influence that neuronal and extra-cranial bodily components have over each other" (2015, p. 3). The EPP approach is supposed to fill in the details here. However, I suspect that this judgement is made a little too quickly, because the dynamical systems description of brain–body–niche interactions is in one sense a higher-level description of those interactions. The dynamical interactions are described as being part of a larger system comprising brain, body, and niche. We can zoom in and focus upon the dynamics of brain or body, but we shouldn't confuse the dynamics of the brain for the dynamics of the overall system. I have highlighted and outlined the neural dynamics required for enculturation in a number of places. For example, in the account of body schema dynamics and in the case of NPP for symbolic cognition, I have outlined the case for dual component transformations (e. g., Menary 2007, pp. 78–83; 2010; 2013 and 2014). Lets take these two cases in order.

In a now famous series of studies, Maravita & Iriki (2004) studied the bimodal interparietal neurons in trained Japanese macaque brains. These neurons respond both to tactile stimulation on the hand (tactile receptive field) and visual stimuli in the same vicinity as the tactile receptive field (the visual receptive field). The visual receptive field was centred on the hand following it through space. When macaques where trained to use a rake to pull food towards them on a table, the observation that struck Maravita and Iriki was that when the macaques used the rake the receptive fields of the bimodal neurons extended along the axis of the rake, including its head. Iriki's interpretation of this is that "either the rake was being assimilated into the image of the hand or, alternatively, the image of the hand was extending to incorporate the tool" (Iriki & Sakura 2008, p. 2230). The extension of the body schema (receptive field) to include the tool happened only during active holding; it reduced to just the hand during inactivity. The interesting result of these experiments is that the existing body schema has the latent capacity to extend to incorporate the tool. LDP can be cashed out in terms of functional changes as the result of scaffolded learning even in the case of macaques, let alone the notoriously plastic brains of humans.

Functional changes can be cashed out in terms of neural redeployment and cortical connectivity. Returning to the case of mathematical cognition, inherited systems for numerosity are evolutionary endowments; we can be reasonably sure of this because they are constant across individuals and cultures and they are shared with other species. The numerosity systems are "quick and dirty"; they are approximate and continuous, not discrete and digital. By contrast, discrete mathematical operations exhibit cultural and individual variation; there is a big difference between Roman numerals and Arabic numerals. They are subject to verbal instruction (they actually depend on language); one must learn to count, whereas one does not learn to subitise. Mathematics depends on cultural norms of reasoning (mathematical norms). The ability to perform exact mathematical calculations depends on the public system of representation and its governing norms. We learn the interpretative practices and manipulative practices as a part of a pattern of practices within a mathematics community, and these practices transform what we can do. They are constitutive of our exact calculative abilities. Mathematical practices get under our skins by transforming the way that our existing neural circuitry functions.

The relationship between the evolutionarily earlier system and the recent development of public mathematical systems, norms, and symbols comes down to the redeployment of the cortical territories that are dedicated to evolutionarily older functions by novel cultural artefacts (e. g., representations, tools). The transformation results in new connections between the frontal lobe for number-word recognition and association, the temporal lobe for the visual recognition of number form, and the parietal lobe for the approximate recognition of magnitudes across both left and right hemispheres (Dehaene 1997).

The deeply transformative power of our learning histories in the cognitive niche relates to the development of our capacities for understanding symbolic representations and for physically manipulating inscriptions in public space. In learning to understand symbols, the first transformation involves our sensorimotor abilities for creating and manipulating inscriptions (the transformation of the body schema). This is something we learn to do on the page and in the context of a learning environment, in public space, before we do it in our heads. Our capacities to think have been transformed, but in this instance they are capacities to manipulate inscriptions in public space.

It looks like PP can provide models of some of the fundamental processing principles at work at the sub-personal neural level, but it is not obvious that it would replace LDP and neural redeployment in the mechanism of transformation. However, Fabry may be right and PP may add another string to the bow of our understanding of how the brain exhibits the plasticity required for cognitive transformation. In that case it provides extra explanatory depth to the account of enculturation, but only as part of a much broader explanatory framework.

3 Some worries for enculturated predictive coding

Fabry provides a persuasive case for how PP could provide the neural underpinnings of enculturation. In this section, however, I will raise some problems for the proposed marriage of CI/ENC and PP. The main issues I will address are as follows:

1. The incompatibility of the isolated brain interpretation (Hohwy 2013) and the active inference interpretation (Clark 2013) of PP.
2. The attempt to explain all cognitive processing in terms of prediction error.
3. The redeployment of neural circuitry as not being explained by PP.
4. The role of NPP as not being explained by the reduction of prediction error.

1. Isolating the brain

If CI/ENC has one central commitment, it is that we should not think of cognition as isolated from the environment. And yet this is exactly how we ought to understand the predictive brain, according to a prominent interpretation of the PP framework. Whenever the PP framework is introduced, it is almost always introduced in the following way: "Accounts of PP generally assume that human perception, action, and cognition are realized by *Bayesian probabilistic generative models* implemented in the human brain. Since *the human brain does not have immediate access to the environmental causes of sensory effects, it has to infer the most probable state of affairs in the environment* giving rise to sensory data" (Fabry 2015, p. 4; my emphasis). The two main motivations for the PP framework are that the brain is isolated from the environment and must make a best guess as to what it is perceiving, and that this kind of probabilistic inference-making results in internal (neurally realized) models of the environment. Putting aside the probabilistic nature of the inferences, this just is old-fashioned individualism. There is a perceptual interface to an environment of hidden variables; the internal system creates internal models (representations) of those hidden environmental variables, which then causally produce behaviour. The internal states must predict the external variables via sensory input, but they have no direct access to the causal ancestry of the sensory input. This form of individualism is used as an explanation for why models and predictions are required: "Because the brain is isolated behind the veil of sensory input, it is then advantageous for it to devise ways of optimizing the information channel from the world to the senses" (Hohwy 2013, p. 238). Hohwy describes the mind–world relation as "fragile" because of the isolation of the brain, and this is why active inference is required.

The saving grace of the PP framework, from the perspective of CI/ENC, is active inference. In Clark's version of PP active inference and cultural props help to minimize prediction errors (Clark 2013); and because of this, there is a deep continuity between mind and world mediated by active inference and the cultural scaffolding of our local niche. Curiously, Hohwy agrees with Clark's interpretation, but at a cost. Hohwy agrees that active inference and the cultural scaffolding of the environment help to change sensory input so as to minimize prediction error, but also "by increasing the precision of the sensory input" (Hohwy 2013, p. 238). According to Hohwy, the primary role of PP is perceptual inference; as a matter of "second order statistics" active inference helps to optimise sensory input so that perceptual inference is less error-prone.

Note the cost. First, active inference and cultural scaffolding is relegated to the secondary role of reducing prediction error for the primary cognitive job of perceptual inference, which is carried out wholly by matching statistical models to sensory input in the brain. Second, Hohwy shows that this interpretation of active inference should be understood against the background of the isolated brain. "The key point I am aiming at here is that this is a picture that accentuates the indirect, skull-bound nature of the prediction error minimization mechanism" (Hohwy 2013, p. 238). Organizing and structuring our environments makes sense if the mind–world relation is fragile in the way that Hohwy presents it, and also because this structuring makes perceptual inference more reliable. I take it that Fabry and Clark would deny this interpretation of the role of active inference and cultural scaffolding. Indeed, Fabry denies Hohwy's 'isolationist' interpretation in her commentary.

However, Fabry does so by playing up the roles of NPP, which go far beyond prediction minimization: "Furthermore, we need to take into account that genuinely human cognitive processes occur in a culturally sculpted cognitive niche. […] These cognitive resources have unique properties that render them particularly useful for the completion of cognitive tasks" (Fabry 2015, p. 12). She also nods to the sub-personal, mechanistic role of PP in the entire brain–body–niche nexus: "[T]he important theoretical contribution made by the prediction error minimization framework is its providing of a sub-personal, mechanistic description of the underlying neuronal and bodily sub-processes" (Fabry 2015, p. 13). It is therefore not clear to me that PP does anything more than provide the functional details of *some* of the neural processing in the brain–body–niche nexus. It certainly should not be taken to provide a comprehensive account of what cognition is and why there is cultural scaffolding, or what its interesting cognitive properties are.[5] It is to these issues that I shall now turn.

2. Everything is predicted

One of the main concerns with the PP approach is that it is used both to try to explain all of cognition and as an explanation of why there is cultural scaffolding. We've already seen a brief hint of this in Hohwy, Clark, and Fabry's work above.[6] The first worry can be found in the expression of PP as originating in the free energy principle:

> The free-energy considered here represents a bound on the surprise inherent in any exchange with the environment, under expectations encoded by its state or configuration. A system can minimise free energy by changing its configuration to change the way it samples the environment, or to change its expectations. These changes correspond to action and perception, respectively, and lead to an adaptive exchange with the environment that is characteristic of biological systems. This treatment implies that the system's state and structure encode an implicit and probabilistic model of the environment. (Friston & Stephan 2007, p. 417)

PP is primarily a model of the way in which top-down processing 'predicts' bottom up sensory input and which samples the environment to change its expectations. These correspond to perception and action respectively[7]. However, it seems odd to build a cognitive theory on the basis of the prediction of sensory signals. This is because much of cognition is not about sensory signal prediction; nor about actions as sampling the environment. Indeed much of cognition isn't about 'prediction' at all. So whilst I agree that at least part of the mechanisms of cognition can be fruitfully modelled by PP, not all of them will be. In enculturated systems, the really important work is being done by the processing governed by normative patterned practices whose properties are understood primarily at the social or populational level. I agree that at the individual level, the mechanisms of ICS can partly be explained by PP, but the main explanatory work will not be a matter of predictions of sensory input[8].

The examples from Landy & Goldstone (2007) may be partly explained by prediction errors, but again this only makes sense in the context of sensorimotor processing governed by mathematical norms. If the norms function as priors in the system, then this might help explain the errors made by the test subjects.

3. Phenotypic plasticity and neural redeployment

PP can't explain the redeployment of neural circuitry to new cognitive functions. And it is not supposed to, since this isn't the job it was designed to do. However, this is a considerable weakness if PP is supposed to be the primary mechanism of enculturation. I've already canvassed the reasons why in section 1.

4. NPP and prediction error minimization

Enculturated PP plays a role in the multilayered and interwoven ICS, but it neither determines nor implements the entire system. My argument in this response has been that the dynamics of ICS are not determined by the predictive processing of parts of the system: if any-

5 CI/ENC provides just these motivations and details. Clark himself proposes that the PP framework "offers a standing invitation to evolutionary, situated, embodied, and distributed approaches to help 'fill in the explanatory gaps' while delivering a schematic but fundamental account of the complex and complementary roles of perception, action, attention, and environmental structuring" (Clark 2013, p. 195).

6 See also their contributions to this volume.

7 There are also theories of attention based upon PP, but I won't address those here.

8 Thomas Metzinger has raised an interesting question for me here: whether there is continuity between the levels? My argument has been that there is continuity between the levels, but this continuity is made possible by NPP's, LDP and neural redeployment. PP explains how we make perceptual inferences about the environment and it might explain something about the hierarchical organisation of neural architecture. However, it should be seen as playing *a* role in the organisation and enculturation of the brain, not the *only* role.

thing PP is enslaved to the processing needs of the entire enculturated system. The PP framework takes perceptual inference as its primary mode of processing, which is the top-down matching of predictions to sensory input. However, it is not obvious that this is the right model for all cognitive processing, since it is not obvious that all cognitive processing is just a matter of predictions about sensory input, nor a hierarchically organised system which minimises prediction errors.

For example Hohwy (2013, p. 238) argues that “many of the ways we interact with the world in technical and cultural aspects can be characterized by attempts to make the link between the sensory input and the causes more precise (or less uncertain).” This would be a very impoverished account of the evolution of public systems of representation. Public systems of representation did not simply evolve to “make the link between the sensory input and the causes more precise (or less uncertain)”; this would be to ignore the social pressures that would have caused representational innovation.[9] It might be true that the history of the refinement of notation has something to do with making input more easy to ‘predict’; however, this would not be an *ultimate* explanation for why there are notations in the first place, nor how they function in our cognitive lives. It *might* be a *proximal* explanation of the neural mechanisms for the processing of notations and as such, it might explain some of the causal conditions that explain how notations have developed, but it doesn’t explain the conditions under which notations evolved. For further reasons why see section 3.4 of my target article, on evolutionary novelty and uniqueness (this volume).

For example, the idea that the brain predicts the product of two numerals makes sense, and the surprise at a product too distant from the operands lends further credence. Remember the example from section 4.1 of my target article (this volume) : 34 + 47 = 268. However, it is not obvious that predictions will help with the second example: 34 x 47 = 1598. What is required in this instance is the serial working through of the multiplication according to an algorithm. Furthermore, this is not simply a case of sensory predictions: when it comes to recognising the numerals on the page in front of you, PP can explain top-down predictions about sensory input, but that is not at all the same thing as the working through of a mathematical problem. So mathematical cognition could not, it seems to me, be reducible to error minimization.

4 Conclusion: Where now?

Despite some of my concerns about how the PP framework can be interpreted and its relation to the CI/ENC framework, I think that Fabry’s account of the enculturation of reading using a hybrid of CI and EPP is really compelling. This leads me to think that an EPP account might by workable for other cases, such as mathematical cognition. Having said this, the division of labour between PP and evolutionary accounts of the origin of NPP and ICS must be in place. The role of scaffolded learning and neural redeployment should not be replaced by error minimization processes. The ‘isolationist’ reading of PP should be resisted, and a more situated cognition friendly approach embraced. PP is a sub-personal account of neural processes that fits within a larger account of the brain–body–niche nexus. If one embraces CI/ENC then there’s more to the mind than What? Now.

Acknowledgements

Thanks to Thomas Metzinger and Jenny Windt for comments and to Regina Fabry for her excellent commentary.

9 I take it that Hohwy is claiming that cultural representations function so as to make perceptual inferences more precise. This would be another way of reducing socio-cultural phenomena to a role that is complementary to the brain, with the processing needs of the brain dictating the evolutionary path that culture must take. The externalist perspective takes it that there are social and cultural pressures that require cognitive innovations (sometimes even new phenotypes).

References

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, *36* (3), 181-204. 10.1017/S0140525X12000477

Dehaene, S. (1997). *The number sense: How the mind creates mathematics.* London, UK: Penguin.

Fabry, R. E. (2015). Enriching the Notion of Enculturation: Cognitive Integration, Predictive Processing, and the Case of Reading Acquisition - A Commentary on Richard Menary. In T. Metzinger & J. M. Windt (Eds.) *Open MIND*. Frankfurt a. M., GER: MIND Group.

Friston, K. & Stephan, K. E. (2007). Free energy and the brain. *Synthese*, *159*, 417-458. 10.1007/s11229-007-9237-y

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

Iriki, A. & Sakura, O. (2008). The neuroscience of primate intellectual evolution: natural selection and passive intentional niche construction. *Philosophical Transactions of the Royal Society B*, *363*, 2229-2241. 10.1098/rstb.2008.2274

Landy, D. & Goldstone, R. L. (2007). How abstract is symbolic thought? *Journal of Experimental Psychology*, *33* (4), 720-733. 0.1037/0278-7393.33.4.720

Maravita, A. & Iriki, A. (2004). Tools for the body (schema). *Trends in Cognitive Sciences*, *8*, 79-86. 10.1016/j.tics.2003.12.008

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

——— (2010). Dimensions of mind. *Phenomenology and the Cognitive Sciences*, *9*, 561-578. 10.1007/s11097-010-9186-7

——— (2013). The enculturated hand. In Z. Radman (Ed.) *The hand, an organ of the mind. What the manual tells the mental* (pp. 593-561). Cambridge, MA: MIT Press.

——— (2014). Neuronal recycling,neural plasticity and niche construction. *Mind and Language*, *29* (3), 286-303. 10.1111/mila.12051

——— (2015). Mathematical Cognition - A Case of Enculturation. In T. Metzinger & J. M. Windt (Eds.) *Open MIND*. Frankfurt a. M., GER: MIND Group.