

# Predictive Processing and Body Representation

Stephen Gadsby<sup>1</sup> and Jakob Hohwy<sup>2</sup>

To appear in: *The Routledge Handbook of Bodily Awareness*. Alsmith, A. J. T. & Longo, M. (eds.). New York: Routledge.

## Abstract

We introduce the predictive processing account of body representation, according to which body representation emerges via a domain-general scheme of (long-term) prediction error minimisation. We contrast this account against one where body representation is underpinned by domain-specific systems, whose exclusive function is to track the body. We illustrate how the predictive processing account offers considerable advantages in explaining various empirical findings, and we draw out some implications for body representation research.

## 1. Predictive Processing and Body Representation

Finding inspiration from a diverse number of fields such as cybernetics, neurobiology, and computational psychiatry, *predictive processing* accounts paint the mind as a neural mechanism whose function is the minimisation of surprise over the long-term average (for introductory statements, see, e.g., Hohwy, 2013; Clark, 2016; Hohwy, 2020). Here, we will present the predictive processing account of body representation, by recounting the framework from the point of view of a single organism, preoccupied with minimising its surprise as best it can.

A biological being, be it a relatively simple organism like *E.coli*, or a more complex organism like a human, exists for a period of time, and then it ceases to exist, at the end of its life. As long as an organism exists, it will tend to occupy only a limited subset of all the possible states it could occupy (for example, humans occupy certain states on land, and rarely in water or in the sky, and *E.coli* tend to occupy states within a certain temperature range). In contrast, when an organism ceases to exist, it will begin to decompose and thereby its remnants will begin to disperse across many different states. This means that, given the kind of organisms we are, it would be *surprising* to see a living human under water or in space, and unsurprising to see it occupy our familiar, land-based, environments. Similarly, it would be surprising to find *E.Coli* in environments much hotter than 70° C. It follows from this trivial observation that organisms, like humans and *E.coli*, must be able to avoid surprising states, or to minimise surprise, given the kind of organisms they are. Otherwise, they simply would not be in existence. Notice that here ‘surprise’ is defined relative to the type of organism, such that what is surprising to one may not be surprising to another (and surprise is a statistical construct unrelated to psychological surprise, such as surprise parties and jump scares). This implies that any organism must somehow be able to form beliefs about which environments are surprising and which are familiar, and then act in the world to avoid the surprising ones and stick to the familiar ones. In a slogan, to exist is to act to minimise surprise.

---

<sup>1</sup> Centre for Philosophical Psychology, Department of Philosophy, University of Antwerp.

<sup>2</sup> Cognition and Philosophy Lab, Department of Philosophy, Monash University.

Predictive processing can now be motivated through this foundational idea of surprise minimisation. The organism needs to know which states are surprising and which are not. The problem is that it cannot know this in advance of encountering any state; this is because the organism cannot not know (in advance) all the states that there are and where it should (most likely) be found among them. In order to minimise surprise, then, it must do something else than directly try to assess surprise. It turns out that if the organism *assumes* a model of what its expected states are, such that any states that differ from that assumed model will appear to be unexpected, or surprising, then it can test predictions generated from this model; doing so will, over some period of continuous trial-and-error testing, help it reduce the real surprise that it cannot directly assess.

The assumed model is a representation of states of the world, and how those states generate particular sensory input. For humans the model would be harboured internally, in the brain; for this reason it is often labelled an *internal generative* model. Technically, this idea is based on developments in physics and machine learning, related to the formal notion of a probabilistic divergence between the assumed model's predictions and the sensory input to the organism, which mathematically must bound the real surprise (for descriptions of the formal background, see Bogacz, 2017; Buckley et al., 2017; Wiese and Metzinger, 2017; Hohwy, 2020).

For present purposes, the important aspect is that as part of this trial-and-error process, the organism can update the elements of its model, in light of the sensory input, with a view to receive less input that contradicts that model's predictions (e.g., if you predict you will have a cup of tea but have run out, then you may revise the model to predict you will have a cup of coffee). Alternatively, it can stick to the model's predictions and seek to plan and execute actions that selectively lead to sensory input that confirms the model's predictions (e.g., if your model makes you expect that you eat every day, and specifies that you tend to eat cake, you may act in the world to make that prediction come true, by baking).

Starting from this simple story about existence, the main components of predictive processing are then lined up. From the assumed model of the expected states of the world, predictions about sensory input can be generated, which can be tested against the actual sensory input. The difference between the prediction and the input is the *prediction error*. The basic idea then is that the organism can help itself to avoid surprising states by keeping track of the prediction error it encounters, given its assumed model. And if the system is going to maintain a good model that rarely encounters surprising states, it should work to keep the prediction error as low as possible, over the long-term average.

Organisms that are overly openminded and change and revise their model in the light of any sensory input they encounter will cease quickly ("I am encountering dangerous tigers, tigers must be part of my expected states"). Similarly, organisms that are overly closeminded will never change their model and will at best selectively sample only the states they expect, and they will also cease to exist quickly, since the world may not cooperate with their prior assumptions ("I expect to win the lottery, so I will spend all my money on lottery tickets"). In other words, organisms that exist for some period of time need to minimise their prediction error, but in a way that *balances* revision of the model, that is, perception and memory, to enable meaningful selective sampling, that is, action. Predictive processing then involves a delicate maintenance of two components: *perceptual inference* of the states of the world and *active inference* of the best plans (or 'policies') for action.

Organisms that perceive and act via perceptual and active inference have to rely on what their assumed model says, that is, on their *prior beliefs* (in the probabilistic sense of ‘beliefs’ as probability distributions or density functions with a mean and a variance). These beliefs give the content of the organism’s representations and in humans are organized in a cortical hierarchy, ordered on the spatiotemporal scale over which the represented states unfold (for example, the visual input as you look out the car window transition over the time scale of seconds, and your states of hunger and thirst transitions over several hours of the day).

Crucially for the topic of body representation, the representational content of the model includes, in addition to the states of the world external to the agent, the agent’s own bodily states. That is, since the body is itself part of what gives rise to sensory input, the agent needs to infer the states of their own body and infer the policies for action that will drive bodily behavior. Here, the predictions generated from these representations can be mixed together, so as to represent interactions among causes that produce relatively abrupt changes in the sensory input (e.g., turning one’s head as one is also sitting in a moving bus, or having trouble getting into a pair of jeans due to both weight gain and shrinkage of the jeans). Being able to conjoin causes and take their interaction into account creates a better fit to the data and less prediction error overall.

Representations in the model will tend to change in light of prediction error, but such revision is subject to maintenance of a balance between being *openminded* and *closedminded*. What determines this balance, and thereby which beliefs are more revisable than others is the *precision* of various beliefs (where precision is the inverse of the variance about the mean of the distribution or density in question). Precision can be thought of as how confident a particular prior belief is, where a very precise belief is more immune to revision than an imprecise belief, because it takes more evidence to overturn the confidence.

Crucially, a predictive processing system also represents and keeps track of the *expected precisions* of the sensory input it comes across or samples, such that its learning rate (i.e., how quickly it changes its representations) can be determined by both how much it already robustly represents and how much it trusts the prediction error associated with some sensory input. The learning rate is also controlled by beliefs about *state transitions* in the environment (sometimes called environmental uncertainty), that is, how likely it is that a particular state may stay the same or change in a more or less predictable or *volatile* fashion. If state transitions are not very probable, then the environment is assumed to not be volatile, and there will be less need to question one’s current representations; in contrast, if the agent has a strong belief that some environment is volatile, then it should relinquish its current representations sooner, and seek out new knowledge.

The internal generative model used in perceptual and active inference provides a whole-of-organism perspective on predictive processing. This perspective implies that there is one, unified model, allowing the agent to operate in its expected states as a unitary, embodied agent, described by the model’s total set of prior beliefs at any given time. The shape of this model is determined by the causal structure of the environments the organism occupies, the internal bodily states of the agent, and the causal interactions between the causes in the external environment and in the body.

On this account, mental representation emerges from what we would call a *domain-general* scheme of predictive processing. While this account explains how organisms come to represent the world (e.g., tea, tigers, and lottery tickets), it also explains how they come to

represent their own bodies (e.g., limb position, body size, touch, glucose levels, heart rate, breath). This is because a generative model whose function it is to model the causes relevant for its sensory input (in order to reduce long-term surprise), inevitably ends up modelling bodily causes too.<sup>1</sup> The idea here is that sensory input caused through the body's interactions with the environment (e.g., the sound of a bottle you dropped on the floor, or the visual input as you turn your head and glimpse your image in a mirror) will be best predicted if the model represents the body as one among the many causes in the world, which combine to produce one's sensory input at any given time (e.g., “the bottle was slippery, my hands were a bit wet, and I am pretty clumsy”; “it was *me* who moved and caused the change in what I saw”). On the predictive processing account, then, body representation simply emerges as the agent attempts to keep minimising surprise as best it can. Of course, it follows that just as we can occasionally be wrong in our perceptual inference about the external world (e.g., when we experience illusions) we can form misrepresentations about our own bodily states.

## 2. Body Representation and Domain Specificity

To minimise surprise about our sensory input, our minds must represent our bodies, because, as described, our bodies are part of what produces our sensory input. Several chapters in this volume are dedicated to characterizing the functional role of the systems that produce these representations. An obvious way of describing their function is as tracking and encoding the body (de Vignemont, 2020a). However, we might ask whether this is the exclusive function of these systems? Put differently: are body representation systems *domain-specific*? We can specify that a system is domain-specific to the extent that “it has a restricted subject matter, that is, the class of objects and properties that it processes information about is circumscribed in a relatively narrow way” (Robbins, 2017). In the context of body representation, this entails that body representation systems are exclusively dedicated to processing information about the body (and perhaps some limited set of associated targets, such as tools that are currently in use or the immediate space surrounding the body). The contrasting view is one where body representations emerge from a domain-general system, responsible for tracking a considerably broader class of targets.

While domain-specificity is a topic of debate in relation to other systems of representation—e.g. language (Pinker, 1994), and social learning (Heyes and Pearce, 2015)—it has not yet received attention in the literature on body representation. Consequently, it is not clear to what extent the domain-specific view is held by researchers. References to systems whose “primary, or fundamental, function” (Goldman, 2012, p. 74) is to represent the body; to “distinct psychological function[s] subserving the orientation on the own body” (Poeck and Orgass, 1971, p. 255); or claims that “the brain evolved [a] dedicated mechanism specifically tuned to the body and its immediate surrounding” (de Vignemont, 2020b, p. 132) appear consistent with (if not positively in favour of) the view. However, our aim is not to dispute any particular account of the function of body representation systems, but to introduce a new way of conceiving of this function—by contrasting it with the domain-specific view. According to our view, body representations emerge from a highly domain-general scheme of prediction error minimisation.

In the next section, we discuss behavioural evidence from two of the most studied domains of body representation: *body size* and *body ownership*. As we illustrate, results from these domains of research suggest that considerably complex inferential resources are brought to bear on the task of body representation. Such resources, we argue, are unlikely to be found in domain-specific body representation systems but are readily available in domain-general systems of environmental modelling. In other words, only if body-representation

emerges in a domain-general representational system would it be predicted that perceptual inference about the body can so seamlessly take onboard multiple, complex, non-body-resources.

### **3. Body Representations and Inferential Complexity**

#### **3.1. Representation of Body Size**

Representations of body size exhibit considerable plasticity: they can be swiftly manipulated by exposing participants to misleading evidence about their own body. An example of this is the Pinocchio illusion. In this illusion, a vibrator is applied to participants' bicep tendon, while they hold on to their nose with the same arm. This change in kinesthetic prediction error (consistent with the bending of the arm) causes an experienced elongation of the nose—as one participant reported: “Oh my gosh, my nose is a foot long! I feel like Pinocchio” (Lackner, 1988, p. 284). It also causes changes in tactile judgment, suggestive of a change in tactile representation of body size (de Vignemont et al., 2005).

Consider two recent variants of this illusion wherein auditory signals are manipulated to achieve the same effect. In one variant, participants were required to pull on the tip of their (occluded) right index finger, while listening to the sounds of rising, falling, or constant pitch (Tajadura-Jiménez et al., 2017). After just two seconds of exposure to the rising pitch sounds (accompanied by finger pulling), participants increased their estimates of finger size and reported an experienced increase in finger size. In the second variant, blindfolded participants were required to drop a ball onto a suspended nylon net (Tajadura-Jiménez et al., 2018). A transducer attached to the net triggered audio-tactile feedback, simulating the impact of the ball hitting the floor. This allowed experimenters to manipulate the delay between the dropping of the ball and the audio-tactile feedback of the ball hitting the floor, simulating the appropriate audio-tactile conditions under which the ball was dropped at taller heights. Following the illusion, participants reported an experienced change in height, increased their estimates of leg length, and exhibited difference in their step size—all indications of a change in represented body size.

These illusions demonstrate that body representations are affected by signals that are not bodily in origin. In each case, the relevant signals were not only associated with non-bodily modalities (i.e. they were auditory, rather than somatic), but, more importantly, they conveyed information about body size via exceedingly indirect inferential routes. In the case of the first experiment, via the association between pitch shifting and motion along the vertical plane; in the second experiment, via expectations of surface collision based on the speed and trajectory of the ball, and the expected distance between its dropping and landing points.

Consider how these results relate to the function of body representation systems. If the exclusive function of such systems is to track the body, it seems unlikely that body representations would be affected by such indirect and abstract inferential routes. In the case of pitch shifting feedback, for example, the relevant signals had no prior association with bodily properties: our bodies do not emit pitch shifting noises, during movement or otherwise. Responding to such information requires monitoring a vastly complex influx of sensory signals and matching these signals against an extensive store of learned associations. If a system was devoted specifically to the domain of body tracking then this would be a remarkably computationally inefficient way of doing so.<sup>ii</sup>

In contrast, the predictive processing account of these illusions is less computationally inefficient. Body representation, on this view, is the upshot of an inference to the best explanation, which draws on all occurrent sensory evidence, as well as domain-general associations between that evidence and more abstract principles regarding environmental regularities. In this case, a crucial piece of the evidence includes the apparently non-accidental simultaneity of the auditory and tactile/kinaesthetic input. Following basic insights about physical causation, the relevant inference is guided by the idea that such constant conjunction between two events indicates causation. Participants then infer the seemingly unique set of beliefs that can make sense of the situation, namely, beliefs that involve a sudden, dramatic, and unusual change in body size. It is unsurprising that such rich and complex inferential processing could emerge from a system whose task is to produce such inferences (and indeed any inferences that are best at minimising prediction error).

### **3.2. Representation of Body Ownership**

Just as we experience the size and shape of our bodies, so too do we experience our bodies as belonging to (or constituting) ourselves. As in the case of representation of body size, careful manipulation of sensory signals can elicit changes in representation of the owned body (see: Ehrsson, this volume). A classic example of this is in the rubber hand illusion (RHI) (Botvinick and Cohen, 1998). In the standard setup of the illusion, participants' (occluded) hands are stroked while, at the same time, they see a rubber hand being stroked (synchronously). Various combinations of such visuo-tactile/proprioceptive input tend to generate an experience that the felt touch is delivered on the rubber hand by the seen stroking tool. The illusion also tends to elicit reports of an experience of ownership towards the rubber hand, along with other responses, such as increased skin conductance when the rubber hand is threatened, suggesting that the rubber hand is represented as belonging to oneself.

Early accounts of the rubber hand illusion posited a two-stage process involving multi-sensory integration of congruent signals and a process wherein a sensory representation of the rubber hand is compared against and/or integrated into a representation of the body (de Vignemont, 2007; Manos Tsakiris, 2010; Carruthers, 2013). In contrast, the predictive processing explanation of this phenomenon appeals in the first instance to precision-weighted multisensory integration, where the expected precision of each sensory input is informed by the context (e.g., the expected high acuity of the visual input in good lighting conditions); in normal conditions this would lead to veridical inferences of tactile sensation and body ownership. But in the rubber hand set-up, the normally well-performing model of the body persistently generates prediction error, due to the mismatch between the felt and seen tactile input. This leads to selection of another model of the body, incorporating the touch on the rubber hand, with a better fit to the incoming sensory input. The deciding factor in selecting between these two models is the strong prior belief, which is independent of the body-specific domain, that synchronous inputs are caused by the things in the environment that are co-located in space and time. On this account, the illusion occurs because, despite its conceptual implausibility, "the rubber hand is mine" model is selected for because, in the situation, it better minimises surprise.

Note that, according to this domain-general process of model selection, there is no limit on the factors that can contribute to a model's superiority. Consequently, a model wherein the rubber hand looked realistic would, all things equal, be a better fit than a model where it did not. This accounts for how cues like the shape and orientation of the hand modulate the

illusion (Tsakiris and Haggard, 2005). However, contrary to early accounts of the illusion, on the predictive processing account, none of these factors are necessary conditions for ownership: while numerous factors contribute to the fit of a model, their absence can be replaced with other forms of evidence (Litwin, 2018). This helps to explain the success of conceptually unusual variants of the illusion, involving ownership towards invisible arms (Guterstam et al., 2013) or non-hand-like objects, such as cardboard boxes (Hohwy and Paton, 2010). Such illusions violate prior beliefs about our bodies and the world, but, in each case, other forms of evidence win out over these inconsistencies. As we will show, more complicated variants of this illusion are also consistent with the domain generality of the predictive processing account.

Several variations of the rubber hand illusion involve the manipulation of non-bodily signals. In one such variation, participants are immersed in a virtual reality air hockey environment and tasked with hitting a puck into a goal (Grechuta et al., 2019a). As in the previous example, this paradigm functions through manipulating audio feedback, creating an experimental condition exhibiting three kinds of incongruency: first, an auditory cue falsely reflecting performance (e.g., a sound of failure following success and vice versa); second, the sound of the hit occurring randomly (within 200–500 ms) before or after the collision; and third, the sound originating from a different location than the collision. This experiment discovered that participants experience a greater sense of ownership over the virtual hand in the congruent, compared to incongruent trials. This difference was evident in self-reported feeling towards the virtual hand, as well as other measures of a change in body representation, including a shift in the proprioceptive representation of hand location (towards the virtual hand) and an increase in galvanic skin response upon threat to the virtual hand.

In a further extension of the paradigm (Grechuta et al., 2019b), the authors manipulated various task-irrelevant environmental features in order to increase incongruence. They manipulated the spatial orientation of audience benches (rotating them on the z-axis); the spatial orientation of a clock (modulating the velocity and direction of the arrows); the position of the sun; and background sounds (to mimic a concert, or cinema). As before, ownership measures were significantly higher in congruent conditions, compared to incongruent.

The predictive processing account of body representation supplies a straightforward explanation for these results: because body representation emerges from a general scheme of environmental modelling, the processes that underpin sense of bodily ownership are expected to be sensitive to expectations about the environment. As the researchers themselves put it, “similar to any robust percept, body ownership depends on the consistency of the internal models of not only the body or the consequences of its actions but also the model of the surrounding environment” (ibid., p. 10). In the case of the incongruent condition, the environment was represented as less trustworthy, so incoming prediction errors were consequently represented with less precision, and therefore were less efficacious in improving the fit of the “the virtual hand is mine” model.

As before, these results appear at odds with the domain-specific view. The aforementioned studies demonstrate that the systems that underpin representation of body ownership are sensitive to highly abstract cues regarding the general reliability of information emerging from an environment. As before, monitoring non-bodily related, abstract features of the environment—such as the functionality of clocks or the spatial orientation of benches—

would be considerably computationally inefficient for a system exclusively dedicated to tracking the body. If, however, expectations about such environmental features are already encoded in the generative model, then those expectations can be swiftly deployed for the task of inferring which feature of that environment belongs to the self.

Finally, consider behavioural results suggesting that the strength of ownership reports in the rubber hand illusion are predicted by participants' hypnotic suggestibility (Lush et al., 2020). In response to these results, some have suggested a different mechanism to explain the characteristic phenomenology of the rubber hand illusion, where top-down 'phenomenological control' through suggestion (via the experimental set-up), rather than a shift in body representation, drives the change in bodily experience. This presents its own domain-general explanation, where the relevant processes do not function to represent the body whatsoever. Such a finding is, however, also consistent with the predictive processing account of bodily ownership, which leaves room for multiple top-down influences on perceptual inference, via the organism's rich web of priors regarding the world. For example, it may be that model selection or prior precisions are modulated when suggestion (hypnotic or otherwise) creates a context of uncertainty. Again, this underscores the flexibility and causal interconnectedness of body representations with other forms of representation.

## **Conclusion**

Predictive processing accounts of cognition are the natural upshot of a necessary condition on existence: the minimisation of surprise. In this chapter, we introduced body representation from this perspective. We illustrated how body representation emerges from a domain-general scheme of prediction error minimisation, contrasting it with the idea of domain-specific, functionally exclusive body representation systems. We illustrated the advantages of predictive processing in terms of explaining behavioural results related to representation of body size and ownership. We finish here by discussing some implications of the proposed view, in terms of researching body representation.

An important upshot of the proposed account is that body representation depends on inferences regarding complex relations between environmental regularities. In some cases, these relations are highly abstract, such as the relationship between pitch shifting and movement trajectories. This invites researchers to look more closely at how body representations can be modified by manipulating signals regarding environmental regularities pertaining to precision and volatility over several time scales and including multiple domains (for examples and review, see, e.g., Limanowski and Friston, 2020; Perrykkad et al., 2021; Smith et al., 2021). Thus, to further our understanding of body representation we should turn our focus away from exclusively manipulating bodily signals, towards integration of more abstract environmental cues (as some researchers have begun to do). In a similar vein, we can understand how misrepresentation of the body occurs—in both pathological and non-pathological contexts—by focusing on how misrepresentation of the environment occurs. An example of this is found in research demonstrating that distortions that were traditionally assumed to be unique to proprioceptive representations of body size (Longo and Haggard, 2010) extend to representation of various environmental objects (specifically, those with which we interact) (Peviani et al., 2021).

There are various further insights to draw from the predictive processing literature, to better understand body representation. For example, there is considerable research into the representational formats employed to model environmental causes, via the surprise



minimisation scheme (Gładziejewski, 2016; Kiefer and Hohwy, 2018). Such insights may help to inform proposals regarding the format and philosophical legitimacy of body representation (Gadsby and Williams, 2018; Alsmith, 2019). Further, there is a rich body of research devoted to understanding representation of the self within a predictive processing framework (Hohwy and Michael, 2017; Letheby and Gerrans, 2017). Such accounts may be used to understand the link between representation of the body and more abstract representations of the self. In turn, we might understand disorders that involve both misrepresentation of the body and disordered representation of the self, such as anorexia nervosa (Gadsby, 2017; Keizer & Engel, this volume), as well as the apparent link between bodily illusions, like the rubber hand illusion, and more abstract beliefs about the self (Krol et al., 2020). However, to reap these explanatory benefits, researchers must first embrace the view that body representation emerges from a predictive processing system that is domain-general rather than domain-specific.

---

<sup>i</sup> The distinction we have drawn here is not between nativism and empiricism. Organisms may begin with some set of very precise homeostatic (i.e. innate) priors (or set points) that specify the states they expect to find themselves (e.g., heart rate, glucose levels, body temperature). Such priors, in virtue of their high precision, may therefore constrain the states in which the body is represented as being in.

<sup>ii</sup> This relates to the claim that domain-specific systems are computationally efficient in virtue of being *informationally encapsulated* (i.e. highly restricted in the class of information that they take as input) (Fodor, 1983).

## References

- Alsmith, A. J. (2019) 'Bodily structure and body representation', *Synthese*, pp. 1-30.
- Bogacz, R. (2017) 'A tutorial on the free-energy framework for modelling perception and learning', *Journal of Mathematical Psychology*, 76, pp. 198-211.
- Botvinick, M. and Cohen, J. (1998) 'Rubber hands 'feel' touch that eyes see', *Nature*, 391(6669) pp. 756-756.
- Buckley, C. L., Kim, C. S., McGregor, S. and Seth, A. K. (2017) 'The free energy principle for action and perception: A mathematical review', *Journal of Mathematical Psychology*, 81, pp. 55-79.
- Carruthers, G. (2013) 'Toward a cognitive model of the sense of embodiment in a (rubber) hand', *Journal of Consciousness Studies*, 20(3-4) pp. 33-60.
- Clark, A. (2016) *Surfing uncertainty: Prediction, action, and the embodied mind*. New York: Oxford University Press.
- Fodor, J. A. (1983) *The modularity of mind*. MIT press.
- Gadsby, S. (2017) 'Distorted body representations in anorexia nervosa', *Consciousness and cognition*, 51 pp. 17-33.
- Gadsby, S. and Williams, D. (2018) 'Action, affordances, and anorexia: body representation and basic cognition', *Synthese*, 195(12) pp. 5297-5317.
- Gładziejewski, P. (2016) 'Predictive coding and representationalism', *Synthese*, 193(2) pp. 559-582.
- Goldman, A. I. (2012) 'A moderate approach to embodied cognitive science', *Review of Philosophy and Psychology*, 3(1) pp. 71-88.
- Grechuta, K., Ulysse, L., Rubio Ballester, B. and Verschure, P. F. (2019a) 'Self beyond the body: action-driven and task-relevant purely distal cues modulate performance and body ownership', *Frontiers in human neuroscience*, 13 p. 91.
- Grechuta, K., de La Torre, J., Ballester, B. R. and Verschure, P. F. (2019b) 'Challenging the boundaries of the physical self: purely distal cues in the environment impact body ownership', *bioRxiv*, p. 672139.
- Guterstam, A., Gentile, G. and Ehrsson, H. H. (2013) 'The invisible hand illusion: multisensory integration leads to the embodiment of a discrete volume of empty space', *Journal of cognitive neuroscience*, 25(7) pp. 1078-1099.
- Heyes, C. and Pearce, J. M. (2015) 'Not-so-social learning strategies', *Proceedings of the Royal Society B: Biological Sciences*, 282(1802) p. 20141709.

- Hohwy, J. (2013) *The Predictive Mind*. Oxford: Oxford University Press.
- Hohwy, J. (2020) 'New directions in predictive processing', *Mind & Language*, 35(2), pp. 209-223.
- Hohwy, J. and Paton, B. (2010) 'Explaining away the body: Experiences of supernaturally caused touch and touch on non-hand objects within the rubber hand illusion', *PloS one*, 5(2) p. e9416.
- Hohwy, J. and Michael, J. (2017) 'Why should any body have a self?' *The subject's matter: Self-consciousness and the body*, 363
- Kiefer, A. and Hohwy, J. (2018) 'Content and misrepresentation in hierarchical generative models', *Synthese*, 195(6) pp. 2387-2415.
- Krol, S. A., Thériault, R., Olson, J. A., Raz, A. and Bartz, J. A. (2020) 'Self-concept clarity and the bodily self: Malleability across modalities', *Personality and Social Psychology Bulletin*, 46(5) pp. 808-820.
- Lackner, J. R. (1988) 'Some proprioceptive influences on the perceptual representation of body shape and orientation', *Brain*, 111(2) pp. 281-297.
- Letheby, C. and Gerrans, P. (2017) 'Self unbound: ego dissolution in psychedelic experience', *Neuroscience of Consciousness*, 2017(1) p. nix016.
- Limanowski, J. and Friston, K. (2020) 'Attentional Modulation of Vision Versus Proprioception During Action', *Cerebral Cortex*, 30(3) pp. 1637-1648.
- Litwin, P. (2018) *Rubber Hand Illusion does not arise from comparisons with internal body models: a new multisensory integration account of the sense of ownership*. PeerJ Preprints. (2167-9843)
- Longo, M. R. and Haggard, P. (2010) 'An implicit body representation underlying human position sense', *Proceedings of the National Academy of Sciences*, 107(26) pp. 11727-11732.
- Lush, P., Botan, V., Scott, R. B., Seth, A. K., Ward, J. and Dienes, Z. (2020) 'Trait phenomenological control predicts experience of mirror synaesthesia and the rubber hand illusion', *Nature Communications*, 11(1), p. 4853.
- Perrykkad, K., Lawson, R. P., Jamadar, S. and Hohwy, J. (2021) 'The effect of uncertainty on prediction error in the action perception loop', *Cognition*, 210, p. 104598.
- Peviani, V., Magnani, F.G., Bottini, G. and Melloni, L. (2021) 'Metric biases in body representation extend to objects', *Cognition*, 206, p.104490.
- Pinker, S. (1994) *The language instinct: How the mind creates language*. Penguin UK.
- Poock, K. and Orgass, B. (1971) 'The concept of the body schema: A critical review and some experimental results', *Cortex*, 7(3) pp. 254-277.

Robbins, P. (2017) 'Modularity of Mind', in Zalta, E. (ed.) *The Stanford Encyclopedia of Philosophy*. Available at: <https://plato.stanford.edu/archives/win2017/entries/modularity-mind> (Accessed: 4 October 2021)

Smith, R., Badcock, P. and Friston, K. J. (2021) 'Recent advances in the application of predictive coding and active inference models within clinical neuroscience', *Psychiatry and Clinical Neurosciences*, 75(1), pp. 3-13.

Tajadura-Jiménez, A., Vakali, M., Fairhurst, M. T., Mandrigin, A., Bianchi-Berthouze, N. and Deroy, O. (2017) 'Contingent sounds change the mental representation of one's finger length', *Scientific reports*, 7(1) pp. 1-11.

Tajadura-Jiménez, A., Deroy, O., Marquardt, T., Bianchi-Berthouze, N., Asai, T., Kimura, T. and Kitagawa, N. (2018) 'Audio-tactile cues from an object's fall change estimates of one's body height', *Plos one*, 13(6) p. e0199354.

Tsakiris, M. (2010) 'My body in the brain: a neurocognitive model of body-ownership', *Neuropsychologia*, 48(3) pp. 703-712.

Tsakiris, M. and Haggard, P. (2005) 'The rubber hand illusion revisited: visuotactile integration and self-attribution', *Journal of experimental psychology. Human perception and performance*, 31(1) p. 80.

de Vignemont, F. (2007) 'Habeas corpus: The sense of ownership of one's own body', *Mind & Language*, 22(4) pp. 427-449.

de Vignemont, F. (2020a) 'Bodily Awareness', in Zalta, E. (ed.) *The Stanford Encyclopedia of Philosophy*. Available at: <https://plato.stanford.edu/archives/fall2020/entries/bodily-awareness> (Accessed: 4 October 2021)

de Vignemont, F. (2020b) 'What Phenomenal Contrast for Bodily Ownership?' *Journal of the American Philosophical Association*, 6(1) pp. 117-137.

de Vignemont, F., Ehrsson, H. H. and Haggard, P. (2005) 'Bodily illusions modulate tactile perception', *Current Biology*, 15(14) pp. 1286-1290.

Wiese, W. and Metzinger, T. K. (2017) 'Vanilla PP for Philosophers: A Primer on Predictive Processing', In Metzinger, T. K. and Wiese, W. (eds.) *Philosophy and Predictive Processing*. Frankfurt am Main: MIND Group

## Further Readings

Fotopoulou, A., (2014) 'Time to get rid of the 'Modular' in neuropsychology: A unified theory of anosognosia as aberrant predictive coding', *Journal of neuropsychology*, 8(1), pp.1-19. (Proposes a unified, non-modular account of anosognosia for hemiplegia, drawing insights from predictive coding and free energy minimisation.)

Limanowski, J. and Blankenburg, F., (2013) 'Minimal self-models and the free energy principle', *Frontiers in human neuroscience*, 7, p.547. (Provides a review and statement of

predictive processing as an integrated account of multisensory integration, interoception, agency, and self.)

Hohwy, J. (2013) *The Predictive Mind*. Oxford: Oxford University Press. (Chapters 5 and 11) (Discusses multisensory integration as a domain-general predictive processing phenomenon and considers the predictive processing aspects of bodily self-awareness.)

Apps, M. and Tsakiris, M., (2014) 'The free-energy self: a predictive coding account of self-recognition', *Neuroscience & Biobehavioral Reviews*, 41, pp. 85-97. (Outlines and defends a predictive coding view of self-representation and uses this view to understand findings from various psychological and neuroimaging experiments.)

Barrett, L.F. and Simmons, W.K., (2015) 'Interoceptive predictions in the brain', *Nature reviews neuroscience*, 16(7), pp.419-429 (Introduces an anatomically inspired predictive coding model of interoceptive processing, with implications for various mental disorders.)