



Attenuating oneself

An active inference perspective on “selfless” experiences

Jakub Limanowski^a (j.limanowski@ucl.ac.uk)

Karl Friston^a (k.friston@ucl.ac.uk)

Abstract

In this paper, we address reports of “selfless” experiences from the perspective of active inference and predictive processing. Our argument builds upon grounding self-modelling in active inference as action planning and precision control within deep generative models – thus establishing a link between computational mechanisms and phenomenal selfhood. We propose that “selfless” experiences can be interpreted as (rare) cases in which normally congruent processes of computational and phenomenal self-modelling diverge in an otherwise conscious system. We discuss two potential mechanisms – within the Bayesian mechanics of active inference – that could lead to such a divergence by attenuating the experience of selfhood: “self-flattening” via reduction in the depth of active inference and “self-attenuation” via reduction of the expected precision of self-evidence.

Keywords

Active inference · Self-model · Sensory attenuation

This article is part of a special issue on “Radical disruptions of self-consciousness”, edited by Thomas Metzinger and Raphaël Millière.

1 Introduction

A key theme of the 2018 “Selfless Minds” workshop¹ was the discussion of reports of “selfless” experiences – conscious episodes during which, if one believes these reports, self-consciousness is partly or even completely lost (e.g., during intoxication or psychosis, cf. Letheby & Gerrans, 2017; Millière, 2017; Saks, 2007). Such

^aUniversity College London

¹The manifesto that guided this workshop has been incorporated into the editorial of this special issue, see Millière & Metzinger (2020).



reports raise an important question for the understanding of selfhood and self-consciousness: How could one have a conscious experience – and able to report on it afterwards – in the absence of any awareness of oneself (as having the experience)?

In this paper, we will approach this question from the perspective of active inference formulations of predictive processing (Friston, 2010; Friston et al., 2010; Friston, Samothrakis, & Montague, 2012). Active inference lends itself to describing brain and mind function, particularly when it comes to the construction of internal models of agents in their lived world (cf. Hohwy, 2013; Clark, 2015; Wiese & Metzinger, 2017). The framework has inspired much conceptual work on the nature of self-modelling and the experience of self (Deane, 2020, this issue; Hohwy & Michael, 2017; Letheby & Gerrans, 2017; Limanowski & Blankenburg, 2013; Metzinger, 2003; Wiese, 2019). A key point of active inference is that predictive processing via internal models underwrites the optimal planning of actions – which rests on the notion of control; i.e., inferring the optimal course of *my* (physical, autonomic, or mental) action² to minimize expected free energy. Based on this formulation of planning (as inference), we will argue that some notion of “selfhood” or “self-agency” – in the sense of inference about control – is inherent in active inference. Crucially, this includes the allocation of precision to sensory evidence, which corresponds to attention as a form of “mental” action (Metzinger, 2017). The problem with reports of “selfless” experiences then boils down to the following question: How come people can *feel like* they lack selfhood, when in fact they are in control of at least some of their behaviour via their self-model? We will argue that these experiences can be interpreted as (rare) cases in which computational and phenomenal self-modelling diverge. We will consider two potential mechanisms – within the Bayesian belief updating of active inference – that could lead to such a divergence by attenuating the experience of selfhood: “self-flattening” via reduction in the depth of active inference and “self-attenuation” via reduction of expected precision of self-evidence.

2 Self-modelling based on predictive processing

First, however, we will briefly introduce some key ideas about internal predictive (self) models in the active inference framework (see Friston, 2010; Hohwy & Michael, 2017; Limanowski & Blankenburg, 2013; Seth & Tsakiris, 2018; Wiese & Metzinger, 2017, for a more exhaustive introduction). Active inference sits within a larger “free-energy principle”, according to which any living system – that can be demarcated from its surroundings – will actively try to remain in a set of unsurprising states by maximizing the (marginal) likelihood of sensory samples (Friston, 2010). In this scheme, free energy minimization corresponds to maximizing

²We use the term “action” here to emphasize the link between active inference in terms of action planning and conscious self-modelling. We will later show how the traditional definition of action as a subset of behaviour, characterized by conscious goal representation and a sense of agency, can be related to the depth of inference.

Bayesian model evidence, which implies a notion of “self-evidencing” (i.e., a Bayes-optimal model – a free energy minimizing agent – will always try to maximize the evidence for its existence, Hohwy, 2016). Such “self-models” are probabilistic (predictive) mappings from causes to consequences, for example from latent or hidden states of the world to sensory observations, in which higher levels contextualize lower levels, and lower levels provide evidence for higher levels (e.g. in the form of “prediction errors”, as in predictive coding; cf. Friston, Rosch, Parr, Price, & Bowman, 2017). This hierarchical scheme of recurrent message passing implies that increasingly higher-level beliefs represent increasingly abstract states of affairs at increasingly broad time scales.

In such deep architectures, balancing the relative influence of prior beliefs or sensory evidence (i.e., prediction errors) on Bayesian belief updating across the entire hierarchy – and between sensory evidence from different modalities – is accomplished by weighting the ascending prediction errors by their relative precision based on prior expectations under the model (Adams et al., 2013a; Feldman & Friston, 2010; Friston, 2010). Precision-modulation is thus also a Bayes-optimal, top-down mechanism that minimizes free-energy by optimally balancing or selecting prediction error signals for hierarchical inference – and implicit Bayesian belief updating – depending on the current context. For instance, prediction errors can be afforded greater precision because they are particularly salient, or because they are particularly relevant for behaviour. A prediction error that is afforded high precision will have a relatively larger impact on inference (i.e. on the updating of the respective prior beliefs). This means that precision has to be estimated and deployed “top-down” at each level of the hierarchy. The functional role of this sort of top-down precision-modulation is equated with attention (Feldman & Friston, 2010; cf. Edwards, Adams, Brown, Pareés, & Friston, 2012). Attention is thus seen as a mechanism by which the impact of sensory evidence on belief updating can be amplified or attenuated (cf. Fazekas & Nanay, 2019), and in this way also accommodates formulations of selective attention, whose allocation is controlled by an interaction of top-down (cognitive) and bottom-up (sensory) factors (Posner, Snyder, & Davidson, 1980; cf. Corbetta & Shulman, 2002; Gilbert & Li, 2013). We will later see how this can be associated with mental action and experienced selfhood.

We (Limanowski & Blankenburg, 2013; Limanowski & Friston, 2018) have previously shown the close correspondence of the formal self-modelling implied by active inference with Metzinger’s (2004) account of *phenomenal* self-modelling; i.e., the construction of a conscious mental model of the organism as a whole (including properties like agency and identity over time). One important assumption of both accounts is that the “self” is seen as a hypothesis or latent state (of being) that can be associated with a self-model. This component of a generative model arises as the (computationally) most accurate, parsimonious explanation for bottom-up multi-sensory information (Metzinger, 2004; cf. Seth, Suzuki, & Critchley, 2011; Allen & Friston, 2016; Apps & Tsakiris, 2014; Ishida, Suzuki, & Grandi, 2015). From a

predictive processing perspective, the hierarchical nature of the underlying computational architecture suggests a centeredness of the model on the “self” (Allen & Friston, 2016; Limanowski & Blankenburg, 2013) in that higher levels of the model will be increasingly abstract (amodal), complex, and invariant (i.e., less likely to be affected by prediction error) – and the highest-level inferred causes pertain to “myself”. We will later see how this is especially important for action planning (i.e., active inference) and, implicitly, for experienced “selfhood”.

Importantly, this (computational) hierarchical notion of self-modelling resonates with the spatiotemporal centeredness of experience on the phenomenal self (cf. Metzinger, 2004), and with the idea of a non-conscious and bodily basis for higher forms of self-consciousness (Blanke & Metzinger, 2009; Gallagher, 2000). Such a framework can therefore be used to explain a vast variety of experimental results and even pathological bodily experience. For instance, bodily illusions are well explained as a result of Bayes-optimal inference; i.e., arising from an interpretation of ambiguous sensory input under strong prior hypotheses (Apps & Tsakiris, 2014; Brown, Adams, Parees, Edwards, & Friston, 2013; Friston, 2005; Limanowski, 2014). In the rubber hand illusion (Botvinick & Cohen, 1998), I “wrongly” adjust my perceived hand position to resolve multisensory ambiguity – but I still feel like a sane person in a normal body with just one, not two right arms (Hohwy, 2013; Limanowski, 2017). Similarly, even when experienced self-location and first-person perspective – two major constituents of minimal phenomenal selfhood along some conceptualizations – are decoupled, a unified self is still experienced (Blanke & Metzinger, 2009; Limanowski, 2014). In a predictive coding scheme, these observations are well explained by the fact that if prediction error can be explained away at lower levels, there is no need to adjust higher-level representations in my model.

In sum, within the predictive processing framework, one can, in principle, associate certain³ computational mechanisms with the phenomenology of “being someone” – in other words, one can link *computational* to *phenomenal* (i.e., conscious mental) self-modelling. But of course, perceptual inference is only part of the self-modelling story. We will next turn to predictive processing as formalized by active inference, which affords a different perspective on self-modelling in terms of action planning and “self-evidencing” (Hohwy, 2016). Specifically, we will discuss how “self-flattening” through reducing the depth of hierarchical (active) inference may play into “selfless” experiences.

³That is, a conscious (phenomenal) self-model is, of course, only a highly specific case of computational “self-modelling”. Of course, there are many instances of computational “self-modelling” that are not accompanied by conscious (self) experience. In this paper, we focus on cases in which processes of phenomenal and computational self-modelling *that are normally congruent* diverge.

3 “Self-flattening”: The relationship between deep active inference and selfhood

Active inference extends perceptual inference or predictive coding by noting action offers another way to quench prediction errors; i.e., sampling sensory data in a way that confirms the model’s predictions. Acting thus involves both generating a prediction of sensory input expected to result from intended movement, and “fulfilling” this prediction by executing the movement, thus effectively suppressing a prediction error signal that would otherwise emerge (Adams et al., 2013a; Brown et al., 2013; Seth & Friston, 2016). The agent must therefore also have beliefs about which course of action (or, more generally, behaviour; see below) will be optimal in a given context. Hence the agent’s model must be able to entertain “counterfactually rich” representations; i.e., beliefs about several alternative potential actions and the states of affairs that these actions would bring about (Friston et al., 2017; Seth, 2014; cf. Powers, 2005). This issue has recently been addressed by a formulation of active inference in terms of Bayesian model selection – among potential courses of action and behaviour – based on their expected free energy (evaluated in the light of prior beliefs and preferences, Parr & Friston, 2017; Friston et al., 2017). In brief, the latent variables of such models⁴ are hidden states and policies; hidden states generate observations and state transitions depend on a plan or action “policy” pursued by the agent. Policy optimization thus entails selecting a sequence of actions, with an associated effect on state transitions and expected outcomes – and a corresponding free energy. In other words, the policy with the lowest *expected* surprise is most probable, i.e., the sort of policy “I am likely to pursue” (Friston, 2018; Friston et al., 2017). As noted above, policies are selected based on inference at multiple levels, where higher levels contextualize lower levels. Thereby, as emphasized in the previous section, active inference relies on sensory evidence and on its appropriate weighting. In turn, the selected policy is one “I am likely to pursue” and will therefore specify empirical “self” priors (for action) that contextualize self-modelling – which again emphasizes the hierarchical nature of self-modelling discussed above (cf. Butz, 2008; Apps & Tsakiris, 2014; Limanowski & Blankenburg, 2013; Seth & Tsakiris, 2018). As has been pointed out by a number of authors, this leads to the notion of “self-evidencing” inherent in active inference (Hohwy, 2016; cf. Friston et al., 2012; Hohwy & Michael, 2017; Limanowski & Blankenburg, 2013): An active inference agent will always try to maximize evidence for the hypothesis it entertains about itself – thus perceptual inference and inferred policies provide this kind of evidence that “I am that sort of agent”.

⁴The generative models in play here are usually conceptualized as discrete state space models similar to Markov decision processes; but the same processes can in principle also be formulated in continuous state space, e.g. in a predictive coding scheme, since both formulations entail (hierarchical) belief updating based on sensory evidence. In discrete state spaces, many inference problems associated with planning and selecting actions become conceptually and mathematically more tractable.

Note that the kind of action (or behaviour in general) we are talking about here is not necessarily physical – the same principles may apply to mental processes (cf. Metzinger, 2017). In particular, one can understand the active deployment of precision as a form of covert or mental action that has exactly the look and feel of attention⁵ (cf. Metzinger, 2004). The argument goes (see Limanowski & Friston, 2018 for details) as follows: Policy optimization necessarily entails a specification of the precision of (action-dependent) changes in hidden states that we are trying to infer. Put in formal terms, policies or beliefs about action entail expectations about precision; placing confidence in the consequences of action is an inherent part of the policies from which we select our actions. This implies that beliefs about actions – in the sense of active inference as action policy optimization – cannot be subject to introspective attention because this would induce another policy (of policies) and an infinite regress. So, these “high level” beliefs about “what I am” doing are unique – and may be the computational basis of precluding an infinite regress by phenomenal self-modelling, with accompanying phenomenal transparency⁶ (Metzinger, 2004; cf. Limanowski & Friston, 2018). This interpretation of active inference speaks to the concept of “attentional agency” as introduced by Metzinger (2013, 2017; cf. Wiese, 2019).

Whether we are talking about mental or physical action, the important point is that policy optimization is a special kind of inference – it is inference about which states of the world I can *control*; i.e., about selecting a course of *my action* that will minimize my expected free energy (cf. Friston et al., 2012). One may now ask: Is this probabilistic representation of control necessarily conscious? An answer to this question has been put forth by Friston (2018) as follows: The representation of action policies – potentially, even several alternative ones, each of which specifies an expectation of how the state of the world (accessible via my sensory states) unfolds depending on my action – requires the system to embody an explicit representation of how states evolve over time. Depending on how far this representation of fictive time (i.e., into the past *and* the future) extends, potential action policies will be temporally deeper. Note that temporal depth is closely related to hierarchical (representational) depth and counterfactual richness (Seth, 2014) because the deeper one goes into the future the greater the number of outcomes. One can now propose an association of temporal depth – the ability to plan and explore multiple futures – with the degree of consciousness⁷ it subtends: whereas non-conscious

⁵This can also be linked to concepts such as introspective attention, defined by Metzinger (2004, p. 36) as a specific kind of introspection (*introspection₁*).

⁶The concept of phenomenal transparency describes the specific case in which only the content of a mental representation, but not its construction process, is available to introspective attention – this may be the reason for why some experiences seem “real” (Metzinger, 2004).

⁷We acknowledge that the idea that consciousness comes in unique “degrees” is not uncontroversial. However, our idea is compatible with multi-dimensional accounts of consciousness (e.g. Bayne, Hohwy, & Owen, 2016): temporal depth would specify the degree of consciousness along a particular dimension – which leaves open whether degrees of consciousness can also be measured along other, perhaps independent dimensions.

processes are stuck in the “here-and-now” (Edelman, 2001), conscious processes operate under a “thick” model of future action and behaviour. This idea speaks to many previous definitions of consciousness as a quintessentially mnemonic process (Damasio, 2012; Edelman, 2003; Husserl, 2008; James, 1890; Seth, 2009; Verschure, 2016; cf. Powers, 2005; Carey, 2018). Note that this sort of temporal depth also grounds the agent in time (it generates a “narrative”, Friston et al., 2017) – and provides an opportunity to explain the often discussed invariance of phenomenal selfhood over time (James, 1890; Metzinger, 2003).

In sum, under active inference, sentient systems that employ temporally thick generative models are likely conscious agents, in that at least some – deep – inference processes are associated with conscious mental states. We believe this idea fits well with the often advanced proposal that there is a basic self-consciousness in the background of any conscious experience (e.g. Zahavi, 2014; Damasio, 1999).⁸ The point we want to make here is that when consciousness arises during (deep) active inference, it will be accompanied by (minimal) consciousness of what one could call the “self-as-agent”. On this view, even in cases of altered self-perception and the absence of overt action, a system may engage in active inference, i.e., in that at least the allocation of attention is controlled (cf. Metzinger, 2013; Wiese, 2019). This would mean that any *conscious* system (including artificial ones) that can be said to engage in active inference will experience some sort of “mental” agency – generating a new transparent layer in its phenomenal self-model. But this seems to contrast with reports of very vivid phenomenology during otherwise “selfless” experiences (Millière, 2017, 2020, this issue; Saks, 2007).

There is an interesting related case of apparent absence of attentional (or more generally, cognitive) control during (relatively) vivid conscious experience: namely, during mind wandering (Metzinger, 2018; Schooler, 2014). Interestingly, by linking the (temporal) depth of inference to consciousness, we can, in principle, accommodate the traditional definition of action as a specific case of behaviour accompanied by a conscious goal representation and sense of agency: there are many kinds of behaviour that do not depend on deep inference (such as homeostasis and reflexes) and are therefore not perceived as (consciously) controlled. Some kinds of behaviour – i.e., actions – are based on deep inference about control and therefore have a phenomenology of agency. This distinction may also apply to different kinds of *attention*; i.e., whereas endogenous attention relies on deep inference and feels “deliberate”, a capture of attention by a salient stimulus feels much less “controlled”. Even though mind wandering episodes may not be characterized by a loss of “selfhood”, they could in principle be linked to a reduced depth of active inference – and may therefore be described as mental behaviour rather than action (Metzinger, 2017).

⁸Whether or not this is to be interpreted as an inherent subjectivity in consciousness, i.e., a pre-reflective self-awareness characterized by the “first-personal mode of givenness” of all conscious experience, is beyond the scope of this paper (Legrand, 2006; Zahavi & Parnas, 1998); but this can be traced back to Sartre, Husserl, and Merleau-Ponty if one wants to).

So, could a potential explanation of “selfless” experiences be that they are related to a reduction of temporal depth of active inference – a “self-flattening” – resulting in an attenuation of phenomenal selfhood (Deane, 2020, this issue)? We think this is unlikely to be the complete explanation, for the following reason. Based on the (computational) argument that inference about the self – and, crucially, which states of the world it can control – is at the core of active inference, we propose that the corresponding processes of inference are also temporally deep(est). Thus any “flattening” of active inference would have general effects on action and perception; i.e., it would involve a general reduction of consciousness as e.g. during certain stages of sleep, anaesthesia, or coma. Moreover, one would expect that if the temporal depth of inference is indeed “flattened” in this way, these experiences should also be less accessible by memory in retrospect. Whether or not this is universally true – and whether or not consciousness is in fact generally reduced during “selfless” experiences – is certainly an empirical question, so our proposal remains speculative.

However, we believe that there is another (not necessarily exclusive) mechanism that could lead (or contribute) to “selfless” experiences; i.e., the control over expected precision. In the next section, we will discuss why it is important to get one’s precision expectations “right” for inference about *what kind of an agent I am*, and how aberrant precision control could play into the sort of self-attenuation that characterises “selfless” experiences, and thus – together with self-flattening – help to explain the apparent differences between computational and phenomenal self-modelling.

4 “Self-attenuation”: The importance of expected precision

As mentioned above, precision control has a fundamental role in the construction of self-representations. Following active inference, this is a general role that should apply to low-level (e.g. bodily) self-representation but – as we will argue – also to higher, cognitive and conceptual levels. Interestingly, the problem that the brain has to solve via precision control is often not which sensory evidence to *emphasize*, but which to *attenuate* (Parr, Rees, & Friston, 2018). This sort of sensory attenuation is especially relevant for action – in fact, it would be impossible to initiate a movement without it.

In brief, on an active inference reading, movement occurs because high-level multi-modal or amodal prior beliefs about behaviour predict proprioceptive and exteroceptive states that would ensue if the movement was performed (e.g. a particular limb trajectory). Prediction error is then suppressed throughout a motor hierarchy ranging from intentions and goals to kinematics to muscle activity (Kilner, Friston, & Frith, 2007). At the lowest level of the hierarchy, spinal reflex arcs suppress proprioceptive prediction errors by “fulfilling” the predicted movement –

which thereby minimises exteroceptive prediction errors (e.g. the predicted visual consequences of the action). The assumption that action is driven by anticipation of its sensory effects links active inference to ideomotor accounts of action (Hommel, Müsseler, Aschersleben, & Prinz, 2001; Prinz, 1997), perceptual control theory (Powers, 2005), and things like the equilibrium point hypothesis for motor control (Feldman, 1974).

Sensory attenuation is also important for the construction of a multisensory body representation – especially when sensory information from multiple modalities is conflicting. In the “rubber hand illusion” (Botvinick & Cohen, 1998), for example, visual information about hand position is expected to be very precise, while the (conflicting) proprioceptive information about hand position is afforded a lower precision – i.e., it is relatively down-weighted – to resolve the intersensory conflict and to maintain a coherent body representation. Similar mechanisms may be in play during visuomotor adaptation, i.e., the adaptation to novel visuomotor mappings as introduced e.g. by wearing prism glasses: some experimental evidence suggests that, as in the rubber hand illusion, a temporary attenuation of irreconcilable proprioceptive information may help with this sort of adaptation to a new body representation (Balslev et al., 2004; Bernier, Burle, Vidal, Hasbroucq, & Blouin, 2009; Limanowski & Friston, 2019). This can potentially go as far as to induce an experiential “neglect” of the real (physical) body during virtual reality experiences – although this has only been shown in monkeys so far (i.e., it has been suggested that monkeys using brain-machine interfaces to control artificial limbs gradually begin to neglect their real body, Carmena et al., 2003; cf. Metzinger, 2007). Lastly, sensory attenuation is crucial for self-other distinction. By attenuating sensory data that is self-produced, I can now emphasize externally generated – behaviourally relevant – sensory data (e.g. during finger movement, self-produced somatosensory input from skin stretching and muscle movements is attenuated, while sensitivity to externally generated touch is enhanced, Limanowski et al., 2019). Thus, it is crucial to know which data to attenuate – this is a problem that any “social” brain has to solve: when interacting with conspecifics, I need to know how to balance proprioceptive and exteroceptive (e.g., visual) information to either move myself, or to be able to observe another’s movements without echopraxia (see Kilner et al., 2007; Friston et al., 2010; Limanowski & Friston, 2019 for further discussion).

Note that getting the precision estimates “right” often means lowering them to attenuate sensory evidence – especially when modelling oneself. In other words – and keeping in mind that the top-down allocation of precision is equated with attention – one can describe this as ignoring or “dis-attending” to certain features of oneself (Clark, 2015; cf. Limanowski, 2014, 2017). In the simplest case, this means lowering the precision afforded to one particular sensory modality (even purely interoceptively, Seth et al., 2011; Allen, Levy, Parr, & Friston, 2019), but we will see that the same principle may hold at more complex levels of self-modelling, too.

The point we want to emphasize is that the temporary attenuation of the precision of sensory “self-evidence” – which is necessary to entertain an alternative (and yet counterfactual, cf. Seth, 2014) hypothesis about myself – is effectively a form of “self-attenuation”. In the case of movement, for instance, self-evidence would be constituted by proprioceptive information – conveying evidence for the fact that I am actually not moving – and attenuation of this self-evidence is necessary to enable movement, i.e., to enact a counterfactual proprioceptive hypothesis issued at higher levels of the hierarchy (Adams et al., 2013a; Brown et al., 2013). A tangible example of sensory attenuation is saccadic suppression, where we appear to be unable to “see” the motion induced by saccadic eye movements. A more fanciful example might be the temporary suspension of attention – which is cued by the misdirection of a magician, but remains under our (top-down precision) control – that allows us to suspend our disbelief that what we are witnessing is indeed “magic”.

Luckily, self-attenuation just needs to be applied transiently – e.g., at movement initiation – and thus sensory evidence can still be processed to guide inference. An interesting thought is that during such periods of “self-attenuation”, self-experience may also be altered. Thomas Metzinger has introduced a related idea with the concept of a “self-representational blink” occurring at the transition to mind wandering (Metzinger, 2013; cf. Wiese, 2019). While this idea may be difficult to test empirically, one could speculate that perhaps some inferred narrative would still maintain the self-model or self-hypothesis during these periods, just like the content of a scene does not disappear during a saccade (cf. Hohwy & Michael, 2017; Seth & Tsakiris, 2018).

What if this sort of attenuation were not temporary? This would, in the long run, have grave consequences for self-experience, action, and ultimately life: Of course, sensory information – self-evidence – is needed to maintain a representation of myself; i.e., the sensory evidence for me as a “self-as-agent”. So, if I continued to attenuate this evidence, my inference about myself would become quite unrealistic and I could not act properly in the world. It is likely that this would ultimately lead to the death of the agent. Less severe examples of how “abnormal” precision control can affect self-experience are found in abundance in psychopathology. On this view, pathological states can be described in terms of a change in the precision of perceptual prediction errors (due to abnormal priors, cf. Friston, 2005; Parr et al., 2018; Sterzer et al., 2018). This theme of aberrant precision control dominates many explanations of false perceptual inference in general and lack of central coherence in psychiatric syndromes in particular, such as organic psychosyndromes (Collerton, Perry, & McKeith, 2005), chronic pain (Tabor & Burr, 2019), functional motor symptoms (Edwards et al., 2012), autism (Lawson, Rees, & Friston, 2014; Pellicano & Burr, 2012; Van de Cruys et al., 2014) and schizophrenia (Adams et al., 2013b; Powers, Mathys, & Corlett, 2017).

In relation to “selfless” experiences, the most interesting pathologies resulting from aberrant precision control are those characterized by *misrepresentations of agency* (Adams et al., 2013b; Brown et al., 2013; Edwards et al., 2012; Limanowski,

2017). For instance, in schizophrenia, inference about the hidden causes of sensations may fail because the precision of high-level beliefs is increased to compensate for a failure to attenuate sensory prediction error during action. These overconfident beliefs generate additional, inappropriately confident, predictions about external causes – the agent is not able to infer whether it caused its sensations itself, or whether someone or something else caused them. This results in altered sensory experience and in severe cases in misattributions of agency – as e.g. demonstrated by different susceptibility of schizophrenic patients in the force-matching paradigm (Brown et al., 2013). The same principles that presumably cause these – and other – hallucinations or delusions (i.e., when the perceptual system is affected, Friston, 2005) can lead to functional symptoms, e.g. when precision control is abnormal in the motor system (Edwards et al., 2012; Parr et al., 2018). In other words, there may be a mechanistic link between dissociative syndromes (known as functional medical syndromes) and reports of selflessness (Edwards et al., 2012). For example, dissociative symptoms such as self-reports of “I cannot feel my arm” are linked to aberrant central processing – as measured with electroencephalography – in empirical studies of sensory attenuation (Hughes, Desantis, & Waszak, 2013; Oestreich et al., 2015; Pareés et al., 2014). Likewise, in functional motor symptoms, abnormal precision seems to be assigned at intermediate levels of the motor hierarchy, which may trigger the execution of a movement – without the accompanying phenomenology of intentional movement generation (which would be associated with higher-level motor areas; Edwards et al., 2012).

To summarize, self-modelling in active inference – and consequently, healthy experience and behaviour – relies on the balance of sensory and prior (model) precision; both abnormally high and abnormally low precision estimates have negative consequences. As we hope to have shown with the above brief review of psychopathology, altered precision expectations (about sensory self-evidence) can affect even the highest levels of the phenomenal self-model – leading to self-other confusion and misattributions of agency. The point we want to make is that the same mechanistic explanation (in terms of sensory attenuation and attentional/precision control) can in principle be extended to account for much of the phenomenology of “selfless” experiences. These experiences could thus be interpreted as (partly) resulting from a temporary attenuation of more complex cues for the self-model (i.e., “self-evidence”), which can apparently be so strong that there is only marginal conscious perception of these cues – leading to a “false” update of the phenomenal self-model. This interpretation particularly speaks to “selfless” experiences of the sort associated with e.g. psychedelics, if one subscribes to the notion that these can be seen as “psychotomimetics” (i.e., that the associated altered state is akin to psychosis, cf. Bayne & Carter, 2018). Such an explanation also aligns with arguments that the sort of “ego-dissolution” reported after some psychedelic experiences may be due to an impairment of those mechanisms that integrate sensory evidence into a coherent self-percept (Letheby & Gerrans, 2017; Millière, 2017).

5 Conclusion and outlook

We have argued that the experience of having “lost” one’s self (constituting a fundamental change to the phenomenal self-model) could arise from a combination of “self-flattening” via a loss of deep active inference and “self-attenuation” via aberrant precision expectations about sensory self-evidence (i.e., within the computational self-model realizing active inference). While the former mechanism could lead to a generally reduced temporal depth and degree of consciousness, the latter could attenuate (aspects of) the “self” from experience. However, even if that system may be “wrong” about what kind of agent it is – including experiences of “selflessness” or self-other confusion – it will still employ an optimal (computational) self-model in the sense implied by the active inference story. We therefore conclude that “selfless” experiences can be interpreted as (rare) cases in which – in an otherwise conscious system – normally congruent processes of computational and phenomenal self-modelling diverge.

This divergence is an exciting area of investigation for interdisciplinary research on self-modelling and -experience. There are many fine details to the active inference story, which we have but touched upon here. For instance, whether or not – and how strongly – the temporal depth of active inference can be “flattened” by e.g. meditation or pharmacology is an exciting empirical question. For example, trained meditators (cf. Berkovich-Ohana, Dor-Ziderman, Glicksohn, & Goldstein, 2013) may be able to direct their attention in a way that will enhance self-attenuation in a very similar way as in pathological or psychedelic experience. In other words, they can evoke changes in neuronal precision control similar to the effects of neuromodulators. This puts precision control at the centre of the (empirical) story again. Furthermore, there are cases in which sensory evidence itself can trigger a loss of transparency – i.e., a revision of beliefs about precision. Such subjectively surprising changes from transparency to opacity encompass, for instance, reaching “lucidity” in a dream; i.e., becoming aware that one is dreaming (Dresler et al., 2015), or certain stress situations; e.g., after accidents, when somehow everything about the situation seems “unreal” (Metzinger, 2003, 2007); they can even be triggered by violations of sensorimotor expectations, such as when an afterimage is recognized as “unreal” because it does not move according to motor predictions sent to the eyes (Seth, 2014). Many altered states of consciousness seem to involve experiences similar to those described above, and could therefore offer an interesting tool to investigate the mechanisms grounding the “realness” of our experience empirically. A detailed investigation of altered states of consciousness under the assumption that they result from aberrant precision-weighting – and, perhaps, an associated loss of phenomenal transparency – could help understand why people may sometimes *feel like* they have “lost” parts of, or even their entire “self”.

Acknowledgments

We thank Thomas Metzinger, Raphaël Millière, and two anonymous reviewers for their helpful comments. This work was supported by funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 749988 to JL. KF was funded by a Wellcome Trust Principal Research Fellowship (Ref: 088130/Z/09/Z).

References

- Adams, R. A., Shipp, S., & Friston, K. J. (2013a). Predictions not commands: Active inference in the motor system. *Brain Structure and Function*, 218(3), 611–643. <https://doi.org/10.1007/s00429-012-0475-5>
- Adams, R. A., Stephan, K. E., Brown, H. R., Frith, C. D., & Friston, K. J. (2013b). The computational anatomy of psychosis. *Frontiers in Psychiatry*, 4(47). <https://doi.org/10.3389/fpsy.2013.00047>
- Allen, M., & Friston, K. J. (2016). From cognitivism to autopoiesis: Towards a computational framework for the embodied mind. *Synthese*, 195, 2459–2482. <https://doi.org/10.1007/s11229-016-1288-5>
- Allen, M., Levy, A., Parr, T., & Friston, K. J. (2019). In the body’s eye: The computational anatomy of interoceptive inference. *BioRxiv*, 603928. <https://doi.org/10.1101/603928>
- Apps, M. A., & Tsakiris, M. (2014). The free-energy self: A predictive coding account of self-recognition. *Neuroscience & Biobehavioral Reviews*, 41, 85–97. <https://doi.org/10.1016/j.neubiorev.2013.01.029>
- Balslev, D., Christensen, L. O., Lee, J. H., Law, I., Paulson, O. B., & Miall, R. C. (2004). Enhanced accuracy in novel mirror drawing after repetitive transcranial magnetic stimulation-induced proprioceptive deafferentation. *Journal of Neuroscience*, 24(43), 9698–9702. <https://doi.org/10.1523/JNEUROSCI.1738-04.2004>
- Bayne, T., & Carter, O. (2018). Dimensions of consciousness and the psychedelic state. *Neuroscience of Consciousness*, 2018(1), niy008. <https://doi.org/10.1093/nc/niy008>
- Bayne, T., Hohwy, J., & Owen, A. M. (2016). Are there levels of consciousness? *Trends in Cognitive Sciences*, 20(6), 405–413. <https://doi.org/10.1016/j.tics.2016.03.009>
- Berkovich-Ohana, A., Dor-Ziderman, Y., Glicksohn, J., & Goldstein, A. (2013). Alterations in the sense of time, space, and body in the mindfulness-trained brain: A neurophenomenologically-guided MEG study. *Frontiers in Psychology*, 4. <https://doi.org/10.3389/fpsyg.2013.00912>
- Bernier, P. M., Burle, B., Vidal, F., Hasbroucq, T., & Blouin, J. (2009). Direct evidence for cortical suppression of somatosensory afferents during visuomotor adaptation. *Cerebral Cortex*, 19(9), 2106–2113. <https://doi.org/10.1093/cercor/bhn233>
- Blanke, O., & Metzinger, T. (2009). Full-body illusions and minimal phenomenal selfhood. *Trends in Cognitive Sciences*, 13(1), 7–13. <https://doi.org/10.1016/j.tics.2008.10.003>
- Botvinick, M., & Cohen, J. (1998). Rubber hands “feel” touch that eyes see. *Nature*, 391(6669), 756. <https://doi.org/10.1038/35784>
- Brown, H., Adams, R. A., Parees, I., Edwards, M., & Friston, K. (2013). Active inference, sensory attenuation and illusions. *Cognitive Processing*, 14, 411–427. <https://doi.org/10.1007/s10339-013-0571-3>
- Butz, M. V. (2008). How and why the brain lays the foundations for a conscious self. *Constructivist Foundations*, 4(1), 1–14. Retrieved from <https://constructivist.info/4/1/001>
- Carey, T. A. (2018). Consciousness as control and controlled perception – a perspective. *Annals of Behavioral Science*, 4(2), 1–8. Retrieved from <http://behaviouralscience.imedpub.com/consciousness-as-control-and-controlled-perception-a-perspective.php?aid=23059>
- Carmena, J. M., Lebedev, M. A., Crist, R. E., O’Doherty, J. E., Santucci, D. M., Dimitrov, D. F., et al. (2003). Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biology*, 1(2), 193–208. <https://doi.org/10.1371/journal.pbio.0000042>
- Clark, A. (2015). *Surfing uncertainty: Prediction, action, and the embodied mind*. Oxford: Oxford University Press.
- Collerton, D., Perry, E., & McKeith, I. (2005). Why people see things that are not there: A novel perception and attention deficit model for recurrent complex visual hallucinations. *Behavioral and Brain Sciences*, 28, 737–757. <https://doi.org/10.1017/S0140525X05000130>
- Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, 3, 201–215. <https://doi.org/10.1038/nrn755>

Limanowski, J., & Friston, K. (2020). Attenuating oneself: An active inference perspective on “selfless” experiences. *Philosophy and the Mind Sciences*, 1(I), 6. <https://doi.org/10.33735/phimisci.2020.I.35>



- Damasio, A. (1999). *The feeling of what happens: Body and emotion in the making of consciousness*. New York: Harcourt Brace.
- Damasio, A. (2012). *Self comes to mind: Constructing the conscious brain*. New York: Pantheon.
- Deane, G. (2020). Dissolving the self: Active inference, psychedelics, and ego-dissolution. *Philosophy and the Mind Sciences*, 1(1), 2. <https://doi.org/10.33735/phimisci.2020.I.39>
- Dresler, M., Wehrle, R., Spoormaker, V. I., Steiger, A., Holsboer, F., Czisch, M., & Hobson, J. A. (2015). Neural correlates of insight in dreaming and psychosis. *Sleep Medicine Reviews*, 20, 92–99. <https://doi.org/10.1016/j.smrv.2014.06.004>
- Edelman, G. (2001). Consciousness: The remembered present. *Annals of the New York Academy of Sciences*, 929, 111–122. <https://doi.org/10.1111/j.1749-6632.2001.tb05711.x>
- Edelman, G. M. (2003). Naturalizing consciousness: A theoretical framework. *Proceedings of the National Academy of Sciences*, 100(9), 5520–5524. <https://doi.org/10.1073/pnas.0931349100>
- Edwards, M. J., Adams, R. A., Brown, H., Pareés, I., & Friston, K. J. (2012). A Bayesian account of “hysteria”. *Brain*, 135(11), 3495–3512. <https://doi.org/10.1093/brain/aws129>
- Fazekas, P., & Nanay, B. (2019). Attention is amplification, not selection. *The British Journal for the Philosophy of Science*, axy065. <https://doi.org/10.1093/bjps/axy065>
- Feldman, A. G. (1974). Change in the length of the muscle as a consequence of a shift in equilibrium in the muscle-load system. *Biophysics*, 19(3), 544–548.
- Feldman, H., & Friston, K. (2010). Attention, uncertainty, and free-energy. *Frontiers in Human Neuroscience*, 4(215). <https://doi.org/10.3389/fnhum.2010.00215>
- Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360(1456), 815–836. <https://doi.org/10.1098/rstb.2005.1622>
- Friston, K. (2018). Am I self-conscious? (Or does self-organization entail self-consciousness?). *Frontiers in Psychology*, 9. <https://doi.org/10.3389/fpsyg.2018.00579>
- Friston, K. J. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11(2), 127–138. Retrieved from <https://www.nature.com/articles/nrn2787>
- Friston, K. J., Daunizeau, J., Kilner, J., & Kiebel, S. J. (2010). Action and behavior: A free-energy formulation. *Biological Cybernetics*, 102(3), 227–260. <https://doi.org/10.1007/s00422-010-0364-z>
- Friston, K. J., Rosch, R., Parr, T., Price, C., & Bowman, H. (2017). Deep temporal models and active inference. *Neuroscience & Biobehavioral Reviews*, 90, 486–501. <https://doi.org/10.1016/j.neubiorev.2018.04.004>
- Friston, K., Samothrakis, S., & Montague, R. (2012). Active inference and agency: Optimal control without cost functions. *Biological Cybernetics*, 106, 523–541. <https://doi.org/10.1007/s00422-012-0512-8>
- Gallagher, S. (2000). Philosophical conceptions of the self: Implications for cognitive science. *Trends in Cognitive Sciences*, 4(1), 14–21. [https://doi.org/10.1016/S1364-6613\(99\)01417-5](https://doi.org/10.1016/S1364-6613(99)01417-5)
- Gilbert, C. D., & Li, W. (2013). Top-down influences on visual processing. *Nature Reviews Neuroscience*, 14, 350–363. <https://doi.org/10.1038/nrn3476>
- Hohwy, J. (2013). *The predictive mind*. Oxford: Oxford University Press.
- Hohwy, J. (2016). The self-evidencing brain. *Nous*, 50(2), 259–285. <https://doi.org/10.1111/nous.12062>
- Hohwy, J., & Michael, J. (2017). Why should any body have a self? In F. de Vignemont & A. Alsmith (Eds.), *The subject's matter: Self-consciousness and the body* (pp. 363–391). Cambridge, MA: MIT Press.
- Hommel, B., Müssele, J., Aschersleben, G., & Prinz, W. (2001). Codes and their vicissitudes. *Behavioral and Brain Sciences*, 24(5), 910–926. <https://doi.org/10.1017/S0140525X01520105>
- Hughes, G., Desantis, A., & Waszak, F. (2013). Mechanisms of intentional binding and sensory attenuation: The role of temporal prediction, temporal control, identity prediction, and motor prediction. *Psychological Bulletin*, 139(1), 133–151. <https://doi.org/10.1037/a0028566>
- Husserl, E. (2008). *On the phenomenology of the consciousness of internal time (1893-1917)* (J. B. Brough, Trans.). New York, NY: Springer.
- Ishida, H., Suzuki, K., & Grandi, L. C. (2015). Predictive coding accounts of shared representations in parieto-insular networks. *Neuropsychologia*, 70, 442–454. <https://doi.org/10.1016/j.neuropsychologia.2014.10.020>
- James, W. (1890). *The principles of psychology*. New York: Holt.
- Kilner, J. M., Friston, K. J., & Frith, C. D. (2007). Predictive coding: An account of the mirror neuron system. *Cognitive Processing*, 8, 159–166. <https://doi.org/10.1007/s10339-007-0170-2>
- Lawson, R. P., Rees, G., & Friston, K. J. (2014). An aberrant precision account of autism. *Frontiers in Human Neuroscience*, 8(302). <https://doi.org/10.3389/fnhum.2014.00302>

Limanowski, J., & Friston, K. (2020). Attenuating oneself: An active inference perspective on “selfless” experiences. *Philosophy and the Mind Sciences*, 1(1), 6. <https://doi.org/10.33735/phimisci.2020.I.35>



- Legrand, D. (2006). The bodily self: The sensori-motor roots of pre-reflective self-consciousness. *Phenomenology and the Cognitive Sciences*, 5(1), 89–118. <https://doi.org/10.1007/s11097-005-9015-6>
- Letheby, C., & Gerrans, P. (2017). Self unbound: Ego dissolution in psychedelic experience. *Neuroscience of Consciousness*, 2017(1), nix016. <https://doi.org/10.1093/nc/nix016>
- Limanowski, J. (2014). What can body ownership illusions tell us about minimal phenomenal selfhood? *Frontiers in Human Neuroscience*, 8. <https://doi.org/10.3389/fnhum.2014.00946>
- Limanowski, J. (2017). (Dis-)attending to the body – action and self-experience in the active inference framework. In T. Metzinger & W. Wiese (Eds.), *Philosophy and predictive processing* (pp. 1–13). Frankfurt am Main, Germany: MIND Group.
- Limanowski, J., & Blankenburg, F. (2013). Minimal self-models and the free energy principle. *Frontiers in Human Neuroscience*, 7. <https://doi.org/10.3389/fnhum.2013.00547>
- Limanowski, J., & Friston, K. (2018). “Seeing the dark”: Grounding phenomenal transparency and opacity in precision estimation for active inference. *Frontiers in Psychology*, 9. <https://doi.org/10.3389/fpsyg.2018.00643>
- Limanowski, J., & Friston, K. (2019). Attentional modulation of vision versus proprioception during action. <https://doi.org/10.1093/cercor/bhz192>
- Limanowski, J., Lopes, P., Keck, J., Baudisch, P., Friston, K., & Blankenburg, F. (2019). Action-dependent processing of touch in the human parietal operculum. *Cerebral Cortex*, bhz111. <https://doi.org/10.1093/cercor/bhz111>
- Metzinger, T. (2003). Phenomenal transparency and cognitive self-reference. *Phenomenology and the Cognitive Science*, 2, 353–393. <https://doi.org/10.1023/B:PHEN.0000007366.42918.eb>
- Metzinger, T. (2004). *Being no one: The self-model theory of subjectivity*. Cambridge, MA: MIT Press.
- Metzinger, T. (2007). Self models. *Scholarpedia*, 2(10). <https://doi.org/10.4249/scholarpedia.4174>
- Metzinger, T. (2013). Why are dreams interesting for philosophers? The example of minimal phenomenal selfhood, plus an agenda for future research. *Frontiers in Psychology*, 4. <https://doi.org/10.3389/fpsyg.2013.00746>
- Metzinger, T. (2017). The problem of mental action – predictive control without sensory sheets. In T. Metzinger & W. Wiese (Eds.), *Philosophy and predictive processing*. Frankfurt am Main: MIND Group.
- Metzinger, T. (2018). Why is mind wandering interesting for philosophers. In K. Fox & K. Christoff (Eds.), *The Oxford handbook of spontaneous thought: Mind-wandering, creativity, and dreaming* (pp. 97–111). New York: Oxford University Press.
- Millière, R. (2017). Looking for the self: Phenomenology, neurophysiology and philosophical significance of drug-induced ego dissolution. *Frontiers in Human Neuroscience*, 11. <https://doi.org/10.3389/fnhum.2017.00245>
- Millière, R. (2020). The varieties of selflessness. *Philosophy and the Mind Sciences*, 1(I), 8. <https://doi.org/10.33735/phimisci.2020.I.48>
- Millière, R., & Metzinger, T. (2020). Editorial introduction. *Philosophy and the Mind Sciences*, 1(I), 1. <https://doi.org/10.33735/phimisci.2020.I.50>
- Oestreich, L. K., Mifsud, N. G., Ford, J. M., Roach, B. J., Mathalon, D. H., & Whitford, T. J. (2015). Subnormal sensory attenuation to self-generated speech in schizotypy: Electrophysiological evidence for a “continuum of psychosis”. *International Journal of Psychophysiology*, 97(2), 131–138. <https://doi.org/10.1016/j.ijpsycho.2015.05.014>
- Pareés, I., Brown, H., Nuruki, A., Adams, R. A., Davare, M., Bhatia, K. P., & Edwards, M. J. (2014). Loss of sensory attenuation in patients with functional (psychogenic) movement disorders. *Brain*, 137(11), 2916–2921. <https://doi.org/10.1093/brain/awu237>
- Parr, T., & Friston, K. J. (2017). Working memory, attention, and salience in active inference. *Scientific Reports*, 7(1), 14678. <https://doi.org/10.1038/s41598-017-15249-0>
- Parr, T., Rees, G., & Friston, K. J. (2018). Computational neuropsychology and Bayesian inference. *Frontiers in Human Neuroscience*, 12(61). <https://doi.org/10.3389/fnhum.2018.00061>
- Pellicano, E., & Burr, D. (2012). When the world becomes “too real”: A Bayesian explanation of autistic perception. *Trends in Cognitive Science*, 16, 504–510. <https://doi.org/10.1016/j.tics.2012.08.009>
- Posner, M. I., Snyder, C. R., & Davidson, B. J. (1980). Attention and the detection of signals. *Journal of Experimental Psychology: General*, 109(2), 160–174. <https://doi.org/10.1037/0096-3445.109.2.160>
- Powers, A. R., Mathys, C., & Corlett, P. R. (2017). Pavlovian conditioning-induced hallucinations result from overweighting of perceptual priors. *Science*, 357, 596–600. <https://doi.org/10.1126/science.aan3458>
- Powers, W. T. (2005). *Behavior: The control of perception*. New Canaan, CT: Benchmark Publications.
- Prinz, W. (1997). Perception and action planning. *European Journal of Cognitive Psychology*, 9(2), 129–154. <https://doi.org/10.1080/713752551>
- Saks, E. R. (2007). *The center cannot hold: My journey through madness*. New York, NY: Hachette Books.

Limanowski, J., & Friston, K. (2020). Attenuating oneself: An active inference perspective on “selfless” experiences. *Philosophy and the Mind Sciences*, 1(I), 6. <https://doi.org/10.33735/phimisci.2020.I.35>



- Schooler, J. (2014). Bridging the objective/subjective divide: Towards a meta-perspective of science and experience. In T. Metzinger & J. Windt (Eds.), *Open mind*. <https://doi.org/10.15502/9783958570405>
- Seth, A. (2009). Explanatory correlates of consciousness: Theoretical and computational challenges. *Cognitive Computation*, 1(1), 50–63. <https://doi.org/10.1007/s12559-009-9007-x>
- Seth, A. K. (2014). A predictive processing theory of sensorimotor contingencies: Explaining the puzzle of perceptual presence and its absence in synesthesia. *Cognitive Neuroscience*, 5(2), 97–118. <https://doi.org/10.1080/17588928.2013.877880>
- Seth, A. K., & Friston, K. J. (2016). Active interoceptive inference and the emotional brain. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1708), 20160007. <https://doi.org/10.1098/rstb.2016.0007>
- Seth, A. K., Suzuki, K., & Critchley, H. D. (2011). An interoceptive predictive coding model of conscious presence. *Frontiers in Psychology*, 2. <https://doi.org/10.3389/fpsyg.2011.00395>
- Seth, A. K., & Tsakiris, M. (2018). Being a beast machine: The somatic basis of selfhood. *Trends in Cognitive Sciences*, 22(11), 969–981. <https://doi.org/10.1016/j.tics.2018.08.008>
- Sterzer, P., Adams, R. A., Fletcher, P., Frith, C., Lawrie, S. M., Muckli, L., & Corlett, P. R. (2018). The predictive coding account of psychosis. *Biological Psychiatry*, 84(9), 634–643. <https://doi.org/10.1016/j.biopsych.2018.05.015>
- Tabor, A., & Burr, C. (2019). Bayesian learning models of pain: A call to action. *Current Opinion in Behavioral Sciences*, 26, 54–61. <https://doi.org/10.1016/j.cobeha.2018.10.006>
- Van de Cruys, S., Evers, K., Van der Hallen, R., Van Eylen, L., Boets, B., de-Wit, L., & Wagemans, J. (2014). Precise minds in uncertain worlds: Predictive coding in autism. *Psychological Review*, 121, 649–675. <https://doi.org/10.1037/a0037665>
- Verschure, P. F. (2016). Synthetic consciousness: The distributed adaptive control perspective. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1701), 20150448. <https://doi.org/10.1098/rstb.2015.0448>
- Wiese, W. (2019). Explaining the enduring intuition of substantiality: The phenomenal self as an abstract 'salience object'. *Journal of Consciousness Studies*, 26(3), 64–87. Retrieved from <https://www.ingentaconnect.com/content/imp/jcs/2019/00000026/f0020003/art00004>
- Wiese, W., & Metzinger, T. (2017). Vanilla PP for philosophers: A primer on predictive processing. In T. Metzinger & W. Wiese (Eds.), *Philosophy and predictive processing*. Frankfurt am Main: MIND Group.
- Zahavi, D. (2014). *Self and other: Exploring subjectivity, empathy, and shame*. Oxford: Oxford University Press.
- Zahavi, D., & Parnas, J. (1998). Phenomenal consciousness and self-awareness: A phenomenological critique of representational theory. *Journal of Consciousness Studies*, 5(5), 687–705. Retrieved from <https://www.ingentaconnect.com/content/imp/jcs/1998/00000005/f0020005/903>

Open Access

This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

