



# Expecting some Action: Predictive Processing and the Construction of Conscious Experience

Kathryn Nave<sup>1</sup> · George Deane<sup>1</sup> · Mark Miller<sup>2,4</sup> · Andy Clark<sup>3</sup>

Accepted: 23 April 2022 / Published online: 10 June 2022  
© The Author(s) 2022

## Abstract

Predictive processing has begun to offer new insights into the nature of conscious experience—but the link is not straightforward. A wide variety of systems may be described as predictive machines, raising the question: what differentiates those for which it makes sense to talk about conscious experience? One possible answer lies in the involvement of a higher-order form of prediction error, termed expected free energy. In this paper we explore under what conditions the minimization of this new quantity might underpin conscious experience. Our suggestion is that the minimisation of Expected Free Energy is not in itself sufficient for the occurrence of conscious experience. Instead, it is relevant only insofar as it helps deliver what Ward et al. (2011) have previously described as a *sense of our own poise over an action space*. Perceptual experience, we will argue, is nothing other than the process that puts current actions in contact with goals and intentions, enabling some creatures to know the space of options that their current situation makes available. This proposal fits with recent work suggesting a deep link between conscious contents and contents computed at an ‘intermediate’ level of processing, apt for controlling action.

## 1 Introduction

Predictive processing offers new insights into the nature, possibility, and structure of conscious experience. By seeing experience as a construct that merges predic-

---

✉ Kathryn Nave  
Kathryn.nave1@gmail.com

<sup>1</sup> University of Edinburgh, 3 Charles St, Edinburgh EH8 9AD, UK

<sup>2</sup> University of Hokkaido, 5 Chome Kita 8 Jonishi, Kita Ward, Sapporo, Hokkaido, Japan

<sup>3</sup> University of Sussex, Sussex House, Falmer, BN1 9RH, Brighton, Japan

<sup>4</sup> Monash Centre for Consciousness and Contemplative Studies, Monash University, Melbourne, Australia

tion and sensory evidence, we begin to see how minds like ours infer the structured world presented to us in experience. This is a perceptual landscape built around two core necessities—the need to select apt world-engaging actions (such as reaching for a glass of water) and the need to maintain the inner milieu within the bounds of human viability (see e.g. Clark (2016), Howhy (2013), Barrett and Simmons (2015)). The two are clearly linked, though the space of apt actions soon outruns the space of actions whose purposes are directly related to keeping us within our species-specific window of biological viability.

Embodied agents are able to explore the space of possible actions by minimizing - see further explanation below - the average subpersonal prediction error expected to accrue over the course of an entire action policy (expected free energy, or EFE). Recent speculations concerning possible links between predictive processing and conscious human experience make essential reference to this quantity, EFE, which is different from standard variational free energy (Friston et al. 2016; Parr and Friston 2019; Millidge et al. 2021). We ask under what conditions the minimization of EFE underpins the emergence of conscious experience. Our suggestion is that minimizing EFE is not sufficient for the presence of conscious experience. Instead, EFE is relevant only insofar as it delivers what (Ward et al. 2011) have previously described - building upon work by Evans (1982) and Grush (2007) - as a sense of our own poise over an action space.

In this paper, we reconstruct the notion of ‘knowing poise over an action space’ as implying an agent-inspectable policy space in which candidate policies afford differing actions. Agent-inspectability is unpacked as fine control over the precision-weighting system enabling an agent to launch and assess multiple simulations of possible futures, so as to optimize contact with their own preferences and intentions. Perceptual experience, we conclude, is nothing other than the process that puts current actions in contact with goals and intentions, enabling some creatures to know the space of options that their current situation makes available. This is different from simply selecting actions that minimise future prediction error. It is being able to step back and inspect a landscape of possible actions. It is the shape of this controllable action-space that, we argue, determines the ‘grain’ of conscious experience—why it is that we perceive cups, shapes, and colours but not lower-level best-guesses such as the location of ‘zero-crossings’. Experience is populated by ‘intermediate level representations’ (Marchi & Hohwy, 2020; Nave 2021) that present actionable properties and features, and what is actionable reflects contingent facts about the creature, its learning history, and its environmental niche.

In section one we give an overview of how active inference can be understood in terms of a formalisation of allostasis, with a focus on how the sense of agency and phenomenal self-modelling is implicit within the active inference framework. In section two we consider what such a broadly applicable story can tell us about the much rarer phenomenon of subjective experience, and consider the limited cases in which there is the development of a hierarchically deep self model and the ability to minimize expected free energy, as a potential marker for the emergence of consciousness. In section three we draw upon the action space account (Ward et al. 2011, which incorporates ideas Grush (2007) and Evans (1982) connecting it (in Sect. 4) to recent work on predictive processing and the spatiotemporal resolution of human action

(Marchi & Hohwy, 2020). In section five we then deepen this connection by showing that free energy minimization grounds the action space story in basic homeostatic imperatives and architectures of control that support a wider repertoire of subsidiary goals and action policies. In section six we pull all these strands together under the banner of ‘generative entanglements’ (Clark 2019) as the mechanistic properties that distinguish the kind of creatures that we call sentient.

## 2 Allostasis and Expected Free Energy

The free energy principle (FEP) redescribes the mutual entrainment of two coupled systems towards a stable attractor in inferential terms, by assigning states that result from random perturbations a free energy, or prediction error, value in terms of how far away (statistically speaking) they take the system away from survival congruent sensory states. The return to typical states is then viewed as formally equivalent to a system successfully minimizing prediction error to infer its most likely state. The term ‘inference’ is a lightweight one, intended to pick out systems with particular dynamical properties, rather than the personal-level notion of reasoning through explicitly encoded propositional beliefs.

As has been argued (Friston 2012, 2013; Ramstead et al. 2018) living systems, as homeostatic ones that survive by maintaining themselves within a viable set of states in the face of environmental perturbation (Cannon 1929), are the paradigmatic case of free energy minimizing dynamics. Interoception—‘sensing the body from within’ (Craig 2002)—is essential for the maintenance of viable internal states. Increasingly, predictive processing mechanisms thought to underlie perceptual inference in exteroception are being applied to understand how the brain performs interoceptive inference (Barrett and Simmons 2015; Seth 2013; Pezzulo 2014). Predictive models on the state of the body allow the system to act to bring internal states—‘essential variables’—into reasonable bounds. For instance, when the brain infers the state of the body to be outside the bounds of a healthy body temperature, it can engage corrective autonomic reflexes—such as perspiring when temperature is too high.

Creatures that track regularities across longer timescales gain the advantage of being able to anticipate dyshomeostatic conditions before they arrive, and act in order to avoid these outcomes. This anticipatory action or ‘predictive regulation’ is known as allostasis (Corcoran and Hohwy 2017; Pezzulo et al. 2015; Schulkin and Sterling 2019; Sterling 2012). Allostasis allows the system to go beyond stimulus driven and reflex based corrective mechanisms. Central to the concept of allostasis is that organisms maintain “stability through change”. Allostasis moves beyond the concept of closed-loop control of homeostatic set points by introducing flexible parameters that can change according to context, such as the anticipatory physiological response to a threatening situation, giving an organism the energetic resources to act to avoid anticipated dyshomeostatic outcomes.

Pezzulo et al. (2015) posit that allostatic action selection is underpinned by hierarchical generative models that apply prior beliefs to map actions or sequences of actions to interoceptive outcomes over time. This mapping of policy-dependent outcomes allows an agent to infer the actions that bring about favourable interocep-

tive, proprioceptive or exteroceptive outcomes—such as realising a state of satiation through a certain sequence of actions. Higher hierarchical levels subtending more temporally deep or extended policy dependent outcomes act to contextualise and guide lower levels. Most fundamentally this is related to realising interoceptive states that are consistent with the organism's survival, referred to as 'self-evidencing' states, the same machinery can be applied such that an organism can not only act on affordances, but can act to bring about future affordances (Pezzulo & Cisek, 2016).

Active inference can then be understood as a formal articulation of allostasis, whereby control problems are cast in terms of a model-based inference about the best action plans (or policies). This is known as *planning as inference* (Botvinick and Toussaint 2012; Kaplan and Friston 2018) under the free energy principle (Friston 2019a). In this scheme, the planning and execution of actions become a problem of inference, where candidate actions are scored with respect to their expected free energy—the average variational free energy the organism expects to accrue in pursuing a given policy. The policy with the least expected free energy is that which is selected. Minimisation of expected free energy concerns both maximising extrinsic (pragmatic) value (defined in terms of prior preferences or goals), and the maximisation of information gain or intrinsic value (reducing uncertainty about the causes of valuable outcomes). This means that in order to minimise expected free energy, the agent must balance fulfilling prior preferences alongside the epistemic goal of reducing uncertainty about the causal structure of the world (Friston et al., 2016; Friston et al. 2015).

Self-modelling mechanisms are central within active inference. For instance, action initiation involves systematic misrepresentation of the state of the self: in moving, the organism predicts itself to be in the sensory state which corresponds to the completed movement (Wiese 2017). This involves *disattending* (Limanowski 2017), that is, lowering precision—about the current sensory evidence about the position of the arm, and attending (allocating high precision) to the desired state (my arm has moved). In other words, all that is required for intentional movement is a specification of the desired sensorimotor endpoint of a movement and motor reflexes bring the motor plant to that equilibrium or setpoint (Adams et al. 2013; Brown et al. 2013; Kilner et al. 2007; Parr et al. 2018). As such, higher level policy predictions are unpacked into lower level motor reflexes, transforming policies into overt actions (Clark 2020; Pezzulo et al., 2015)

The system also attenuates sensory evidence of the sensory consequences of self-generated movements. In classic motor control, this was understood in terms of a 'comparator model', whereby the system compares expected and actual consequences of action. Monitoring the mismatch between actual and potential consequences of actions was originally argued to underpin refinement of motor commands, and has since been invoked to understand the origin of the sense of agency (Hohwy and Michael 2017; Lukitsch 2020). In active inference, this is simply part of the top-down prediction about the sensory consequences of actions, but the principle remains that in monitoring the ongoing (mis)match between actions and outcomes, the system can infer its endogenous control over sensation via action.

In selecting actions and action policies with the least expected free energy, an organism has to infer *itself* as able to bring about the consequences of the action

— as an endogenous controller of sensation. While perception in active inference can be understood in terms of inference about hidden states of the causes of sensory signals, action selection requires inference about transitions between states contingent on actions. Importantly, this inference about how sensation is controlled via action tracks both proximal action consequences; for example the sensorimotor contingencies involved in turning over a tomato, and distal ones; the abstract future consequences associated with a given action or action policy, where the degree of abstraction increases with temporal depth (Pezzulo et al. 2018). This inference about control, inherent in active inference, has been associated with the phenomenology of being a *self* or an agent in the world (Hohwy and Michael 2017; Limanowski and Friston 2018, 2020), and has been dubbed ‘agentive control’ (Deane 2020; Deane et al. 2020).

### **3 Birds Do It, Bees Do It, Even Homeostatic Machines Do It, But What Makes a Free Energy Minimizer Conscious?**

From steam engines to stock traders (automated or otherwise) regulatory systems are everywhere. Anticipatory regulatory systems may be less common, but we can plausibly count colonies of *Burkholderia proteobacteria* (Goo et al. 2012), jumping spiders (Schomaker 2004), and self-driving cars among their number. One of the central offerings of the active inference formalization of allostasis—and of cybernetic formalizations of homeostasis that preceded it (Ashby 1952)—is the means of abstracting beyond the organism, to identify these same dynamics across a wide range of physical systems. In this spirit, the Free Energy Principle has been applied not only to brains and bacteria, but also to coupled pendulums (Kirchhoff et al., 2018), the Watt governor (Baltieri, Buckley & Bruineberg, 2020), and an oil drop suspended in water both in support, and as criticism, of Friston’s assertion that it stands as a theory of “every ‘thing’ that can be distinguished from other things in a statistical sense” (Friston 2019a).

So what does active inference have to say about the distribution of a conscious experience across this diverse array of systems and scales? To take the dynamics of active inference as being the hallmark of a conscious being would result in an extremely liberal view of what suffices for the attribution of sentience indeed. Conservatism about consciousness aside, we know from our own lack of experience that the majority of our regulatory processes, from shivering when cold to going through a morning routine on autopilot, need not involve conscious awareness.

One possibility is to argue that active inference and predictive processing in themselves are simply not about consciousness at all. In taking this route Anil Seth and Jakob Hohwy (2021) propose that it is precisely because predictive processing is not itself a theory of consciousness, that it provides an ideal foundation for building such a theory. The emerging consensus in active inference is that consciousness is grounded in the self-modelling processes that are inherent in the action selection and precision control mechanisms in active inference. In this way, consciousness emerges only in systems that minimize expected free energy by the selection of action policies over time. However, this may still provide at most a necessary condition on conscious

experience. As Friston (2018) points out, not all active inferencers are equal. Within the general class of free energy minimizing systems we can identify those with a more temporally-thick model, granting the capability to infer far into the future; with the counterfactual-depth needed to mentally explore the consequences of possible non-actual actions, and the development of a self-model granting the capability to involve one's own projected future needs into that calculation. Such capacities, he proposes, are just what is needed to sort those that are conscious, from those that are not:

“One could then describe systems that have evolved thick generative models (with deep temporal structure) as agents. It now seems more plausible to label these sorts of systems (agents) as conscious, because they have beliefs about what it is like to act; i.e., just be an agent. Furthermore, because active inference is necessarily system-centric the self-evidencing of motile creatures can only be elevated to self-consciousness if, and only if, they model the consequences of their actions. Put simply, this suggests that viruses are not conscious; even if they respond adaptively from the point of view of a selective process. Vegans, on the other hand, with deep (temporally thick) generative models are self-evidencing in a prospective and purposeful way, where agency and self become an inherent part of action selection” (ibid).

Consciousness, then, emerges in systems that evaluate deep and counterfactual rich models of the world. These requirements of temporal depth and counterfactual thickness of a generative self-model do seem to track the behavioural capacities that mark the transition to increasingly sophisticated forms of life—the difference between an *E. coli* bacterium's ability counteract a drop in glucose levels by breaking down glycogen, versus my ability to prepare a sandwich in anticipation of becoming hungry later this afternoon. They do not, however, provide a clear cut answer to whether something becomes conscious or not—only a graduated scale against which a particular system might be assessed. As Friston et al. (2020) acknowledge the view of consciousness as emerging from increasing hierarchical modelling layers, “entails that there is only a gradual difference between some non-conscious and conscious systems, and that consciousness is a vague concept.”

But suppose we are willing to accept this—willing to consider the question “conscious or not” as one that admits of degree. Having a temporally-thick, counterfactually deep, self-model seems an obvious prerequisite for the presence of higher-order self-consciousness, but it is not immediately clear why it would be necessary for the emergence of more basic phenomenal properties—such as the immediate visual experience of an apple on the table. In order to explain why we believe that it is then, we need to step away from the FEP, and PP for a moment, to take another look at the content of such simple visual experiences.

#### 4 The Action-Space Account of Perceptual Content

We know from our own case that even in creatures like ourselves, visual experience is typically not required for the co-ordination of a surprisingly sophisticated reper-

toire of behaviours. There is a wealth of empirical evidence that visual awareness is unnecessary to visually-guide behaviors, such as pointing, tracking and reaching—all of which ordinary participants are able to perform without conscious perception of their target (Bridgeman et al., 1981; Goodale et al., 1986; Castiello et al. 1991).

This becomes particularly striking in neurological disorders such as visual agnosia and action-blindsight (Dancerkert and Rossetti 2005) where damage to areas of the brain involved in visual processing significantly disrupts phenomenology without equivalent impairments to visual action-guidance. For instance, the well-studied visual-agnosic, DF (Whitwell et al. 2014; Milner and Goodale 2006) who lacks visual awareness of the size or shape of a slot in front of her eyes, yet, when instructed to do so, can post a letter through that same slot with perfect ease.<sup>1</sup>

So conscious visual perception may be significantly impacted without impairing online visual guidance of pre-established regulatory-routines—unconscious vision seems to handle this well enough on its own. This, as Clark (2001, 2007, 2009) notes, poses a problem for accounts that seek to account for the content of perceptual experience solely in terms of how it serves the guidance of ongoing action. Instead Ward et al. (2011), develop the ‘action space’ account of perceptual experience—building on previous work that attempts to account for the content of perceptual experience in terms of how it relates to an individual’s dispositions towards an array of goal-directed behaviours, and possession of the visuomotor skill to coordinate these Evans 1982; Grush 2007—see also Schellenberg 2007; Briscoe 2014). As they propose:

“... what counts for (what both explains and suffices for) visual perceptual experience is an agent’s direct unmediated knowledge concerning the ways in which she is currently poised (or, more accurately, the way she implicitly takes herself to be poised) over an ‘action space.’” (Ward, Roberts & Clark, 2011 p.383)

The absence of this direct awareness of the range of action-routines currently available to her is manifest in DF’s behavioural capacities. To talk of her impairment as merely perceptual, as Milner and Goodale (2006) initially did when using DF’s case to support the division of visual processing into unconscious ‘vision for action’ and conscious ‘vision for perception,’ misleadingly implies that the capacity for action is entirely unaffected by the loss of conscious perception. Such a strong division overlooks not only the fact that DF’s inability to produce utterances appropriate to incoming visual stimulation is itself a kind of action, but also that she is further unable to: indicate the width of the slot with her hands; match the orientation of a letter to that of the slot without posting it; take the initiative to post the letter without prompting; or to scale her grasp to a briefly presented object after any delay.

DF can use visual input to guide the ongoing unfolding of a pre-specified action towards a target. She can also initiate actions within a known (e.g. dining) context, within which preserved available visual cues - for example, the metallic sheen of the

---

<sup>1</sup> Our argument does not rest on the claim that D.F lacks visual experience altogether—The point is only that her visual experience changed radically upon the damage to her ventral stream and that such a difference in experience is not ‘purely perceptual’, but is rather inseparable from consequent impairments to action-planning capacities.



cutlery - can serve to guide and structure the action. What she lacks, Ward, Roberts, and Clark argue, is the ability to know, in new and unfamiliar situations, just how she is poised over an action space. That's why, for example, she is unable to immediately choose the right grip for the functional use of an unfamiliarly shaped object, or a familiar one when shown without an identifying context - see Carey et al. (1996), and further discussion in Goodale and Milner (2004), Clark (2009). In the Carey et al. (1996) experiment there is no supporting context, and all objects were selected to be grey, to minimise colour-cues. Under those conditions, DF would often fail to select a functionally apt grip. In this broad way, her visual processing impairments prevent her from automatically putting visual information in touch with the kinds of high-level planning that would deliver better results - on this, see Milner and Goodale's (2006) distinction between simple action 'programming' and higher-level, more intuitively 'abstract' planning.

The upshot is that in many unfamiliar situations DF would exhibit a lack of sensitivity to changing environmental affordances, depriving her (we would argue) of 'immediate knowledge of her own poise over an action space.' PP/Active Inference, we argue, provides a means to operationalize this (in DF compromised) form of online action planning by framing cognitive operations as involving in the attempt to minimize error across multiple timescales, via a hierarchical predictive architecture. Simple online motor control in a familiar context amounts to the sending down of a cascading prediction that guides action to bring one's incoming sensory signal in-line. Fully detached, non-visual reasoning corresponds to revising higher levels of the model to reduce internal incongruities—absent the transmission of any action-eliciting predictions, and consequent feedback from, the sensorimotor periphery. When an "agent's perceptual sensitivity is such as to automatically mesh with her capacities for intentional activity" (Ward et al. 2011) we have a circular feedback loop spanning higher and lower levels, with longer term predictions constraining the operations of the lower levels, while continuous/high-precision error signals at the lower levels seep up to trigger adjustments at these higher levels in turn. This enables new higher level plans to be created 'on the fly' that mesh with specific, changing worldly opportunities.

In the case of DF, we suggest that it is this flagging of worldly opportunities that has in some way broken down. Although she can, if prompted, engage the letterbox with a posting action, her visual encounters (thanks to the damage to her ventral stream) simply do not present her with a rich realm of flagged (salient) possibilities for action.

## 5 Giving the Action Space an Allostatic Foundation

Like the preceding picture, and older cousins such as disposition theories (Evans 1982; Grush 2007; Briscoe 2014) the action space view ties our visual experience to our capacities for bodily action. The crucial difference between this and other action-based accounts of perception, such as the sensorimotor theory (Noë & O'Regan, 2001, 2003), or actionism (Noë 2010), is that the relevant skill does not concern knowledge of the precise patterns of sensory stimulation our movements will bring



about. Instead the content of our action space is constructed at a level more coarse-grained than exact sensorimotor trajectories, and more selectively organized around our particular goals and interests. As Ward, Roberts & Clark. write: “An action space, in this specific sense, is to be understood not as a fine-grained matrix of possibilities for bodily movement, but as a matrix of possibilities for pursuing and accomplishing one’s intentional actions, goals and projects” (Ward et al., p 383).

In moving to explain perceptual content in terms of the content of a space of intended actions, it lacked, however, an account of how these goals and the structure of the action spaces develops, of the principles guiding this selectivity and coarse-graining. As such it appears to fall foul of what Hurley (2001, 1998) termed, ‘the myth of the giving’—that is, of attempting to explain perceptual content in terms of ‘just more content’ as though the content of our intentions could be taken as explanatorily primitive. Marchi and Hohwy’ make progress here in arguing that the content of our basic actions will be constrained by the limitations of our capability to exert control over particularly fast-changing, or abstract, predictions. Yet this cannot be the whole story. As Nave (2021) points out inattentional blindness (Simons & Chabris, 1999), gist perception (Alvarez and Oliva 2009) and global precedence effects (Navon 1977) suggest that the granularity of our perceptual experience is variable and often neglects relatively stable and high-level occurrences (such as the presence of a gorilla) that we should, in principle, be capable of acting upon. Nor is the maximal resolution of our experience fixed by hardwired constraints on basic action alone, as Li et al. (2009) have demonstrated in showing how repeated engagement in actions that require making unusually fine-grained discriminations—such as playing video games like Call of Duty—can lead to a dramatic improvement in a person’s contrast sensitivity.

So, there must be further factors, beyond mere *capability* for control that determine the shifting, and context-dependent, scope of perceptual experience. In the remainder of this paper, we argue that adequately addressing Hurley’s myth of the giving requires a further story about *desirability* of control. That is an explanation of why our intentions, our action space, and so our perceptual experience, is weighted towards the control of some things and not others. This, we argue, is exactly what active inference and layered control bring to the table, by providing an account of how the construction of higher-level action spaces, may be founded upon a basic and fundamental imperative to keep our allostatic self within viable bounds.

Beginning with the phylogenetically wired-in prediction that my blood sugar level should be 70 to 100 mg/dL, I can also learn the regularity of this dropping (prediction error increasing) as 1pm approaches. Moreover, I can learn that not only does releasing glucose stores reduce this prediction error, but also that eating a meal at 12pm can prevent this prediction error from arising at all. Thus I begin to predict daily meal-eating at 12pm. If I’ve not started eating as midday approaches, then prediction-error now results. This increases the salience of, previously irrelevant, action possibilities which I have learnt will reduce this prediction error—the trajectory towards the opening of the fridge, removing the tub of soup and placing it in the microwave. On the view of perceptual experience we are developing, learning this regularity changes the contents of our visual experience, that is the known landscape of salient possibilities for the pursuit of our goals, even in cases where the set of things that are potentially

visible to the eye is unchanged. That what we see depends upon our skills and interests is supported by empirical evidence, such as inattentional blindness (Simons and Chabris 1999) and visual sensitivity training (Li et al. 2009).

New trajectories are not all we learn, however. We may also progressively discover at what level of detail an action trajectory needs to be enforced at in order to result in successful control. A wide array of specific sensorimotor trajectories, corresponding to opening the fridge, will interchangeably minimize prediction error relative to the expectation of the cup being in our hands. Thus, when it comes to exploring potential control strategies, we learn that we do not need to individually explore each and every possible combination of sensorimotor signals. Instead, these details can be condensed into a single course-grained action possibility, that can then be evaluated for its potential to bring our sugar level back within expected bounds. Which trajectories stand out to us for exploration, and at which level of detail, will shift not only with changes in the external world, but also as our internal situation changes. When I'm anticipating a drop in blood sugar levels, it is the pathways to the fridge, the soup, and the microwave that dominate. When I'm tired, my exploration of potential actions skews towards the sofa, the remote control and the television.

This gives us a picture of how the action space account can be augmented by understanding the self in terms of a model of allostatic control model (Deane 2020, 2021; Deane et al. 2020). The agentic control in the allostatic control model is not an impassive observation of the control of sensory inputs via action. Rather, it is concerned with bringing the sensorium in line with the set of states that an organism has learnt to expect as those consistent with its own survival. To do this, as we have noted, involves selecting policies associated with the least expected free energy, where expected free energy can be decomposed into the pragmatic and epistemic value associated with the given policy (Friston et al. 2015). This means, in selecting a policy an organism seeks both to fulfil prior preference (for example, become satiated), and also to resolve uncertainty about the world. This bridges the current proposal to accounts grounding selfhood in the body (Allen & Friston, 2018; Apps and Tsakiris 2014; Limanowski and Blankenburg 2013; Seth 2014).

By tracking how well it is faring in bringing about self-evidencing outcomes, the organism can infer the precision on its own action model—that is, an inference about control in a given context. Precision on the action model can be understood as an inference about control or realisation of prior preferences. endogenous ability to control prior preferences. Affective inference tracks context specific precision on (for instance) prior preferences. For example, violation of the “healthy body condition” expectation (Ongaro and Kaptchuk 2019) manifests to the system as pain, and the system must act to bring sensations into line with the expectation (where the expectation here is physiological integrity). Bayesian and predictive coding mechanisms pointing to the fact that pain seems to be inferential (Anchisi and Zanon 2015). Painful percepts, on this view, integrate prior beliefs with the current sensory evidence, weighted by their expected precision in the context (Morton et al. 2010). Incorporating contextual factors into affective inference means affective responses can be ‘tuned’ such that their motivational salience tracks the needs and goals of the organism. Stress-induced analgesia is one illustration of this—the pain of a twisted ankle should not be motivationally salient when trying to outrun a bear.

High-level priors in the generative model can track temporally deep outcomes, and as such, failure to meet an expected *rate* of prediction error reduction over time manifests to the system as negative affective (Joffily and Coricelli 2013; Kiverstein et al., 2019; Van de Cruys 2017; Hesp et al. 2021). A better than expected rate of prediction error over time manifests to the system as positive affect (consider unexpected rewards, either those that fulfil prior preferences like ice cream, or unexpected epistemic rewards—an “aha!” moment for instance).

The upshot of this is that what is salient to an organism is not constrained to, for instance, the smell of food when hungry. Instead, organisms with hierarchically deep contextualization of interoceptive signals are tuned to appropriate action and engagement with environmental affordances (Pezzulo & Cisek, 2016), and assign appropriate weight to priors and ascending prediction errors across the cortical hierarchy, including context dependent gain control in sensory cortices. Even the smell of hot food may not be salient to an organism engaged in, for instance, finishing a paper ahead of a deadline.

In this way, precision on policies inference about *intentional selection* (‘what am I doing’) — determines *attentional selection* (What am I seeing?). Attentional orientation—where precision is assigned in the cortical hierarchy, depends on hierarchically deep interoceptive inference on the situation. The control context is similarly vital—inference on what “I can” (Bruineberg 2017) do determines where attention should shift to update the model to perform allostatic action and behavioural control, according to goals. Inferences about action-outcome contingencies (control), across multiple timescales, informs how precision (attention) is assigned accordingly relative to the organisms inferred control.

This proposal thus delivers on both the action-oriented and the affective dimensions of consciousness. As Ullman et al. put it, ‘We implicitly but continually reason about the stability, strength, friction, and weight of the objects around us, to predict how things might move, sag, push, and tumble as we act on them.’ (2017, p. 649). At the same time, integrating these expectations of control with inference about (temporally deep) expectations of self-evidencing outcomes, furnishes this proposal with the affective dimensions of consciousness (Deane 2021), such that we encounter “a structured world apt for action and intervention, and inflected at every level, by an interoceptively-mediated sense of mattering, reflecting ‘how things are for me as an embodied agent’” (Clark 2019, p7).

## 6 Intermediate Level Processing

Within the circular feedback loops needed to bridge low-level visual information and higher-level abstract planning, predictions that depict objects and properties at a certain grain or level play a special role. Marchi and Hohwy (2020) and Nave (2021) address what the former refer to as the ‘scope question’ viz. ‘what we are conscious of, given that we are conscious at all’. They make a compelling case for the claim that ‘intermediate level representations’ play a special role in determining the scope of conscious contents within an active inference framework. These representations are constructed within the predictive hierarchy at a level that is neither too abstract

nor too fine-grained to guide the selection of policies for basic action. According to this picture, it is the spatiotemporal resolution of typical human actions that determines this level, which might thus vary for different organisms. The neural realizers of conscious contents, they argue, is determined by the role they play in selecting actionable policies.<sup>2</sup>

Their proposals build on previous ideas concerning the privileged status of certain intermediate-level representations, beginning with Jackendoff (1987) and continuing through Prinz (2000), Koch (2004), and Prinz (2012). The general idea is that intermediate level representations sit between characterizations that are too abstract to determine one action rather than another, and those that are too low-level. As Nave (2021) puts it

*“Our actions are coarse grained in their objects—operating on averages and indifferent to the fine degrees of variance that our perceptual systems are in principle capable of registering. We do not act on photons, or on the thumbnail-sized orientated bars detected in the early visual cortex. We act on a landscape of graspable units, movable objects, and stable surfaces. Small variations in low level details alter this landscape not a jot.”*

.Thus, when we encounter a world of actionable objects and states of affairs, we do not experience every low-level nuance in our own processing. For example, we do not experience the computations of ‘zero-crossings’ that seem to underlie edge and boundary detection, or the multiple low-level hypotheses that must be varied and updated as we look at an object from various angles. Instead, all we see is a bound, shaped object, rotating in egocentric space. Nor do we visually experience highly abstract properties such as pure object-hood.

Strikingly, the higher and lower bounds of the experiential realm seem to reflect the kinds of information suited to the selection of one kind of basic action over another. The idea then, is that some levels of the generative model preferentially depict a possibility space for organismically basic action. It is only at those privileged levels that precise actionable policies can be inferred.

This proposal is a good fit, we suggest, with the action space account. To succeed in our intentional plans and projects, we need to infer and implement actionable policies. The phenomenal realm, if their proposal is on track, is constructed so as to enable us to encounter our world in a way that is ready-parsed for the kinds of basic action we can perform. Precise policies (in the active inference sense of ‘precise’) are ones that will deliver high amounts of reliable prediction error, so that minimizing those errors controls fluent successful action. Neither the very low-level nor the very highest-level control states (the ones that determine tiny bodily nuances or drive long-term projects and goals such as writing a book) fit this bill. Longer-term policies such as writing a book are precise and actionable only to the extent that they are

<sup>2</sup> There are other explanations as to why the intermediate level might be privileged that could be advanced from an efficient coding perspective (see Fleming and Storrs 2019; Green 2019). Thanks to an anonymous reviewer for pointing this out.

composed of, or reliably give rise to, sequences of policies that engage basic actions in this way.

Optimal performance demands that the selection of local action policies is consistent with the availability of precise control. It is this demand that explains the privileged status of the intermediate level information that seems to populate phenomenal experience. The scope of the phenomenal realm, they argue, is delimited by the level of the generative model at which the space of actionable policies can be safely explored. Flagging that level stops us from constantly inferring policies that we cannot implement. As Marchi and Hohwy put it: “a conscious agent is conscious of the hypotheses that are flagged at the appropriate resolution for optimal inference of policies allowing efficient and successful performance of sequences of basic actions (control states)” (p. 17).

## 7 Generative Entanglements

Pulling these various strands together yields a concrete proposal concerning the machinery responsible for conscious experience. Deeply entangled with our grip on the outside world, an inward-looking (interoceptive) cycle targets our own changing physiological states—states involving the gut, viscera, blood-pressure, heart-rate, and the whole inner economy underlying hunger, thirst, and other bodily needs and appetites. As our bodily state alters, the salience of various worldly opportunities (to eat, for example) alters too. That means I will also act differently, harvesting different streams of information. Philosophers and psychologists talk here of ‘affordances’ (Bruineberg and Rietveld 2014), where these are the opportunities for action that arise when a certain type of creature encounters a certain kind of situation—for example, a hungry green sea turtle encountering a nice patch of algae discovers an affordance for eating. But the sea turtle that has just eaten may not find the next patch of algae quite so attractive. Such creatures orient towards the changing value of different affordances given their changing bodily needs.

That already captures a form of very basic sentience. We can think of basically sentient beings as those whose neural model of the (organism-salient) world is in constant two-way communication with their own changing physiological state. Such creatures will perceptually encounter a world fit for action, in which what actions are selected depends heavily upon their current and on-going bodily state and needs. But this falls short, we have argued, of delineating the conditions responsible for true conscious awareness.

What’s missing is something just a little bit ‘higher order’. The creature we just imagined is in touch with its world, in a way that brings together bodily (allostatic) needs and the opportunities for action made available by the sensed environment. But to truly experience that world, the information available to drive action needs to be in some elusive sense ‘available to the creature in question’. We have tried to unpack this notion by suggesting that the creature needs not simply to act, but should find itself confronted with an action space—a perceptual array that affords multiple responses, and that (in so doing) is in touch with capacities for planning and intentional action.

This inserts a kind of gap between sensory stimulation and action, one that sometimes results in characteristic behaviors such as pausing to ponder what to do next. It may be that the sea-turtle finds itself poised over just such a space. It seems extremely unlikely that the bacterium does so, even though it too displays allostatic responses and integrates bodily and worldly information as a means of determining next actions.

The gap, we suggested, is nothing other than a set of opportunities for control. It reflects the availability for control of action of information computed using a temporally deep generative model. Such a model needs to integrate bodily and exteroceptive information with goals and purposes at various timescales. When this occurs, there is the possibility for conflict between possible policies and courses of action. To be poised over an action space is thus to become informed of potentially conflicting possibilities in a way that invites further attempts at optimization—for example, stopping to reflect on what we are about to do.

Pezzulo et al. (2018) note the important role of motivation in this process. We are not simply poised over an action space, nor are we simply poised over an action space that is allostatically inflected. Instead, we are poised over an action space in a way that makes contact with a complex, multi timescale set of motivations and priorities. To borrow their example, finding myself in a restaurant confronted with the dessert trolley, I encounter an action space ( a space of affordances) that is brought into contact with my own longer-term goals and wishes. According to their picture, a ‘control hierarchy’ (plausibly associated with the activities of dorsolateral PFC) then exchanges messages with a ‘motivational hierarchy’ (plausibly associated with activities in ventromedial PFC), so as to drive action selection in a way sensitive to long-term goals and wishes, such as the wish to avoid sugary desserts—items that may be presenting undeniable short-term allostatic attractions. Motivations, reflecting both immediate context and longer term goals, alter the weightings (precisions) assigned to opportunities revealed by the control hierarchy. Motivated action occurs when high precision is assigned to one of the opportunities for action. At that moment, we are driven to act on our knowledge of what we can do. In this way, the revealed action space is placed in direct contact with affect, goals, and motivation, operating over temporally extended periods. The ‘feel’ of being poised to act reflects this combination of knowledge about control (what we can do) and knowledge about motivations and goals. In DF, the downstream impairment to areas crucial for visual form recognition restricts the kinds of information available for this kind of integration.

A particularly attractive element of this story is that it generalises to poise over mental action space as well. Mental action in the active inference framework builds straightforwardly on the planning as inference story. In this case, however, the inference about hidden states concerns attentional states rather than the hidden states causing sensory impressions, and the state transitions refer to transitions between attentional states (Smith et al. 2020). Recall, in active inference, policy selection involves selecting the sequence of actions associated with the least expected free energy, based on beliefs about transitions between states. Inference about the hidden states themselves—perceptual inference or ‘state estimation’—is based on a likelihood mapping that encodes beliefs about how the (hidden) states in the world relate to the observations they generate. Attentional processes are understood in terms of the ‘precision’— the second order confidence in this likelihood mapping. In other

words, the precision can be understood as an estimate of confidence that the agent has that their observations reliably map to hidden states in the world. In the same way that the agent can infer control over sensory inputs via action, and select policies accordingly, the (implicitly) metacognitive modelling of attentional states means that the agent can perform covert (mental) actions in order to perform state transitions between attentional states.

If this complex multi-dimensional story is on track, then experience emerges where (i) there is integrated bodily and worldly information computed using a generative model that displays temporal depth, and (ii) where that model integrates control and motivation across many timescales, bringing goals and affect into direct contact with an appreciation of the space of possible actions that are currently enabled. When those twin conditions (resulting in a highly complex set of ‘generative entanglements’—see Clark 2019) are met, a creature knows the value-inflected action space that is currently made available by its own perceptual contact with the world.

## 8 Conclusions

What our account aims to capture is the process of optimally inferring genuinely actionable policies, where that requires exploring possibilities defined at the intermediate level of the generative model. Creatures with very limited action repertoires have no need to engage mental exploratory mechanisms of this kind. Nor do creatures whose reactions to stimuli are all reflex-like, hence defined very close to the sensory stimulations themselves. But creatures whose generative models span many spatiotemporal scales, and whose projects and goals take many forms, are creatures that might otherwise find themselves constantly inferring policies that they cannot successfully implement. Creatures like that will benefit from a privileged mode of representation that presents the world as a kind of inspectable domain populated by possibilities for action and intervention.

Marchi and Hohwy (2020) and Nave (2021) stopped short of claiming that intermediate-level flagging is identical with phenomenal consciousness, arguing only that it resolves the ‘scope’ question—the question of what contents get to populate conscious awareness. We suggest that locating their considerations within the larger organizing framework of both allostasis and the action-space account makes it plausible to assert a slightly stronger claim. Much of the lived character of perceptual conscious experience, we suggest, reflects the process we have been exploring—the means by which intermediate-level contents become poised for the control of intentional action.

As part of that process, interoceptive predictions turn a bare space of action-possibilities into an allostatically sensitive arena: a space populated by opportunities to serve bedrock organismic needs. This is populated by predictions at the intermediate level of the generative model because this is the level of detail at which our allostatic goals and means of achieving them are specified. Too low a level, and we end up selecting between irrelevant distinctions, such as the precise trajectory with which I will grasp my glass of water. Too high a level, and we do not pick out specific actionable objects, but rather classes or sets of objects. In complex creatures with



temporally deep models, the saliency map of our action space may become increasingly detached from current bodily needs. By supporting a complex and multi-level notion of self-preservation across timescales this temporally deep model allows that the weighting given to an action which increases sugar intake may be sublimated to the desire to maintain a healthy diet which may in turn be weighed out by the desire to maintain my friendship with the person who has baked this cake.

Poised over these enriched action spaces, we are immersed in a world of reasons and possible actions. Does this now explain consciousness itself? Probably not. What is on offer is just a kind of engineering blueprint for a being that would (if it was able to be interrogated) say that it encounters a space of opportunities for action and one whose selected actions would suggest (to an outside observer) that it cares about how it negotiates that space. Believers in a ‘hard puzzle’ of conscious experience will regard this as leaving all the difficult work still to be done. But hard problem sceptics (such as Dennett (1991) (2015), Frankish (2019)) may find in our picture additional resources to explain some of the patterns of behavior and response most typically associated with the apparent presence of that special something - phenomenal consciousness itself.

**Conflict of interest** None.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Adams, R. A., S. Shipp, and K. J. Friston. 2013. Predictions not commands: active inference in the motor system. *Brain Structure and Function* 218 (3): 611–643.
- Allen, M., and K. J. Friston. 2018. From cognitivism to autopoiesis: towards a computational framework for the embodied mind. *Synthese* 195 (6): 2459–2482.
- Alvarez, G. A., and A. Oliva. 2009. Spatial ensemble statistics are efficient codes that can be represented with reduced attention. *Proc. Natl. Acad. Sci. U.S.A.* 106, 7345–7350.
- Anchisi, D., and M. Zanon. 2015. A Bayesian perspective on sensory and cognitive integration in pain perception and placebo analgesia. *PloS one* 10 (2): e0117270: 1–20
- Apps, M. A., and M. Tsakiris. 2014. The free-energy self: a predictive coding account of self-recognition. *Neuroscience & Biobehavioral Reviews* 41: 85–97.
- Ashby, W. R. 1952. *Design for a brain*. London, UK: Chapman and Hall.
- Baltieri, M., C. L. Buckley, and J. Bruineberg. 2020, July. Predictions in the eye of the beholder: an active inference account of Watt governors. In *Artificial Life Conference Proceedings* (pp. 121–129). One Rogers Street, Cambridge, MA 02142—1209 USA journals-info@mit.edu: MIT Press.
- Barrett, L. F., and W. K. Simmons. 2015. Interoceptive predictions in the brain. *Nature reviews neuroscience* 16 (7): 419–429.
- Botvinick, M., and M. Toussaint. 2012. Planning as inference. *Trends in cognitive sciences* 16 (10): 485–488.

- Bridgeman, B., M. Kirch, and A. Sperling. 1981. Segregation of cognitive and motor aspects of visual function using induced motion. *Perception & Psychophysics* 29 (4): 336–342.
- Briscoe, R. 2014. Spatial content and motoric significance. *AVANT. Pismo Awangardy Filozoficzno-Naukowej* 2: 199–217.
- Brown, H., R. A. Adams, I. Parees, M. Edwards, and K. Friston. 2013. Active inference, sensory attenuation and illusions. *Cognitive Processing* 14 (4): 411–427. <https://doi.org/10.1007/s10339-013-0571-3>.
- Bruineberg, J. 2017. Active Inference and the Primacy of the ‘I Can’. In *Philosophy and Predictive Processing: 5*, eds. T. Metzinger, and W. Wiese, Frankfurt am Main: MIND Group. doi:10.155027/9783958573062.
- Bruineberg, J., and E. Rietveld. 2014. Self-organization, free energy minimization, and optimal grip on a field of affordances. *Frontiers in human neuroscience* 8: 599.
- Cannon, W. B. 1929. Organization for physiological homeostasis. *Physiological reviews* 9 (3): 399–431.
- Carey, D. P., M. Harvey, and A. D. Milner. 1996. Visuomotor sensitivity for shape and orientation in a patient with visual form agnosia. *Neuropsychologia* 34: 329–338.
- Castiello, U., Y. Paulignan, and M. Jeannerod. 1991. Temporal dissociation of motor responses and subjective awareness: A study in normal subjects. *Brain* 114 (6): 2639–2655.
- Clark, A. 2001. Visual experience and motor action: Are the bonds too tight? *Philosophical review* 110 (4): 495–519.
- Clark, A. 2007. What reaching teaches: Consciousness, control, and the inner zombie. *The British Journal for the Philosophy of Science* 58 (3): 563–594.
- Clark, A. 2009. Perception, action, and experience: Unraveling the golden braid. *Neuropsychologia* 47 (6): 1460–1468.
- Clark, A. 2016. *Surfing Uncertainty*. NY: Oxford University Press.
- Clark, A. 2019. Consciousness as generative entanglement. *The Journal of Philosophy* 116 (12): 645–662.
- Clark, A. 2020. Beyond desire? Agency, choice, and the predictive mind. *Australasian Journal of Philosophy* 98 (1): 1–15.
- Corcoran, A. W., and J. Hohwy. 2017. *Allotaxis, interoception, and the free energy principle*. Feeling our way forward.
- Craig, A. D. 2002. How do you feel? Interoception: the sense of the physiological condition of the body. *Nature reviews neuroscience* 3 (8): 655–666.
- Danckert, J., and Y. Rossetti. 2005. Blindsight in action: what can the different sub-types of blindsight tell us about the control of visually guided actions? *Neuroscience & Biobehavioral Reviews* 29 (7): 1035–1046.
- Deane, G. 2020. Dissolving the self. *Philosophy and the Mind Sciences* 1 (I): 1–27.
- Deane, G. 2021. Consciousness in active inference: Deep self-models, other minds, and the challenge of psychedelic-induced ego-dissolution. *Neuroscience of Consciousness* 2021 (2): niab024.
- Deane, G., M. D. Miller, and S. Wilkinson. 2020. Losing Ourselves: Active Inference, Depersonalization and Meditation. *Frontiers in Psychology* 11: 2893.
- Dennett, D. 2015. Why and How Does Consciousness Seem the Way it Seems? In *Open MIND*, eds. Thomas Metzinger, and Jennifer M. Windt, Frankfurt am Main: MIND Group. 1–11
- Dennett, D. 1991. *Consciousness Explained*. Boston: Little Brown.
- Evans, G. 1982. *The Varieties of Reference*. Oxford University Press.
- Fleming, R. W., and K. R. Storrs. 2019. Learning to See Stuff. *Current Opinion in Behavioral Sciences* 30: 100–108.
- Frankish, K. 2019. The Meta-Problem is The Problem of Consciousness. *Journal of Consciousness Studies* 26 (9–10): 83–94.
- Friston, K. 2012. Prediction, perception and agency. *International Journal of Psychophysiology* 83 (2): 248–252.
- Friston, K. 2013. Life as we know it. *Journal of the Royal Society Interface* 10 (86): 20130475.
- Friston, K. 2018. Am I self-conscious?(Or does self-organization entail self-consciousness?). *Frontiers in psychology* 9: 579.
- Friston, K. 2019a. A free energy principle for a particular physics. *arXiv preprint arXiv:1906.10184*.
- Friston, K., F. Rigoli, D. Ognibene, C. Mathys, T. Fitzgerald, and G. Pezzulo. 2015. Active inference and epistemic value. *Cognitive neuroscience* 6 (4): 187–214.
- Friston, K., T. FitzGerald, F. Rigoli, P. Schwartenbeck, and G. Pezzulo. 2016. Active inference and learning. *Neuroscience & Biobehavioral Reviews* 68: 862–879.
- Friston, K. J., W. Wiese, and J. A. Hobson. 2020. Sentience and the origins of consciousness: From Cartesian duality to Markovian monism. *Entropy* 22 (5): 516.

- Goodale, M., and D. Milner. 2004. *Sight unseen: An exploration of conscious and unconscious vision*. Oxford: Oxford University Press.
- Green, E. J. 2019. "A Theory of Perceptual Objects." *Philosophy and Phenomenological Research* 99 (3): 663–693.
- Grush, R. 2007. Skill theory v2. 0: Dispositions, emulation, and spatial perception. *Synthese* 159 (3): 389–416.
- Goo, E., C. D. Majerczyk, J. H. An, J. R. Chandler, Y. S. Seo, H. Ham, ... I. Hwang. 2012. Bacterial quorum sensing, cooperativity, and anticipation of stationary-phase stress. *Proceedings of the National Academy of Sciences*, 109(48), 19775–19780.
- Goodale, M. A., D. Pelisson, and C. Prablanc. 1986. Large adjustments in visually guided reaching do not depend on vision of the hand or perception of target displacement. *Nature* 320 (6064): 748–750.
- Hohwy, J. 2013. *The Predictive Mind*. NY: Oxford University Press.
- Hohwy, J., and J. Michael. 2017. 16 Why should any body have a self?. *The subject's matter: Self-consciousness and the body*, 363.
- Hurley, S. L. 1998. *Consciousness in action*. Harvard University Press.
- Hurley, S. 2001. Perception and action: Alternative views. *Synthese* 129 (1): 3–40.
- Joffily, M., and G. Coricelli. 2013. Emotional valence and the free-energy principle. *PLoS Comput Biol* 9 (6): e1003094.
- Kaplan, R., and K. J. Friston. 2018. Planning and navigation as active inference. *Biological cybernetics* 112 (4): 323–343.
- Kilner, J. M., K. J. Friston, and C. D. Frith. 2007. Predictive coding: an account of the mirror neuron system. *Cognitive processing* 8 (3): 159–166.
- Kirchhoff, M., T. Parr, E. Palacios, K. Friston, and J. Kiverstein. 2018. The Markov blankets of life: autonomy, active inference and the free energy principle. *Journal of The royal society interface* 15 (138): 20170792.
- Kiverstein, J., M. Miller, and E. Rietveld. 2019. The feeling of grip: novelty, error dynamics, and the predictive brain. *Synthese* 196 (7): 2847–2869.
- Li, R., U. Polat, W. Makous, and D. Bavelier. 2009. Enhancing the contrast sensitivity function through action video game training. *Nature neuroscience* 12 (5): 549–551.
- Limanowski, J. 2017. (Dis-) attending to the Body: Action and Self-experience in the Active Inference Framework.
- Limanowski, J., and F. Blankenburg. 2013. Minimal self-models and the free energy principle. *Frontiers in human neuroscience* 7: 547.
- Limanowski, J., and K. Friston. 2018. 'Seeing the dark': grounding phenomenal transparency and opacity in precision estimation for active inference. *Frontiers in psychology* 9: 643.
- Limanowski, J., and K. Friston. 2020. Attenuating oneself. *Philosophy and the Mind Sciences* 1 (1): 1–16.
- Lukitsch, O. 2020. Effort, Uncertainty, and the Sense of Agency. *Review of Philosophy and Psychology* 11: 955–975.
- Millidge, B., A. Tschantz, and C. L. Buckley. 2021. Whence the expected free energy? *Neural Computation* 33 (2): 447–482.
- Milner, D., and M. Goodale. 2006. *The visual brain in action* (Vol. 27). OUP Oxford.
- Morton, D. L., W. El-Deredy, A. Watson, and A. K. Jones. 2010. Placebo analgesia as a case of a cognitive style driven by prior expectation. *Brain research* 1359: 137–141.
- Nave, K. 2021. *Visual experience in the predictive brain is univocal, but indeterminate*. 1–25. Phenomenology and the Cognitive Sciences.
- Navon, D. 1977. Forest before trees: The precedence of global features in visual perception. *Cognitive psychology* 9 (3): 353–383.
- Noë, A. 2010. Vision without representation. Perception, action, and consciousness: sensorimotor dynamics and two visual systems, 245–256.
- Noe, A., and J. K. O'Regan. 2003. On comprehending the sensory effects of movement: Toward a theory of perception and consciousness. *Perception ECVF abstract* 32: 0–0.
- Ongaro, G., and T. J. Kaptchuk. 2019. Symptom perception, placebo effects, and the Bayesian brain. *Pain* 160 (1): 1.
- Parr, T., and K. J. Friston. 2019. Generalised free energy and active inference. *Biological cybernetics* 113 (5): 495–513.
- Parr, T., G. Rees, and K. J. Friston. 2018. Computational neuropsychology and Bayesian inference. *Frontiers in human neuroscience* 12: 61.

- Pezzulo, G. 2014. Why do you fear the bogeyman? An embodied predictive coding model of perceptual inference. *Cognitive, Affective, & Behavioral Neuroscience* 14 (3): 902–911.
- Pezzulo, G., F. Rigoli, and K. Friston. 2015. Active Inference, homeostatic regulation and adaptive behavioural control. *Progress in neurobiology* 134: 17–35.
- Pezzulo, G., and P. Cisek. 2016. Navigating the affordance landscape: feedback control as a process model of behavior and cognition. *Trends in cognitive sciences* 20 (6): 414–424.
- Pezzulo, G., F. Rigoli, and K. J. Friston. 2018. Hierarchical active inference: a theory of motivated control. *Trends in cognitive sciences* 22 (4): 294–306.
- Ramstead, M. J. D., P. B. Badcock, and K. J. Friston. 2018. Answering Schrödinger's question: A free-energy formulation. *Physics of life reviews* 24: 1–16.
- Schellenberg, S. 2007. Action and self-location in perception. *Mind* 116 (463): 603–632.
- Schomaker, L. 2004, October. Anticipation in cybernetic systems: A case against mindless anti-representationalism. In *2004 IEEE International Conference on Systems, Man and Cybernetics (IEEE Cat. No. 04CH37583)* (Vol. 2, pp. 2037–2045). IEEE.
- Schulkin, J., and P. Sterling. 2019. Allostasis: A brain-centered, predictive mode of physiological regulation. *Trends in Neurosciences* 42 (10): 740–752.
- Schwarz, N. 1990. *Feelings as information*. Informational and motivational functions of affective states. eds. E. T. Higgins, and R. M. Sorrentino, *Handbook of motivation and cognition: Foundations of social behavior* (Vol. 2, pp. 527 ± 561). New York: Guilford Press
- Seth, A. K. 2013. Interoceptive inference, emotion, and the embodied self. *Trends in cognitive sciences* 17 (11): 565–573.
- Seth, A. K. 2014. *The cybernetic Bayesian brain*. Frankfurt am Main: MIND Group: Open MIND.
- Anil K Seth & Jakob Hohwy. 2021. Predictive processing as an empirical theory for consciousness science. *Cognitive Neuroscience* 12:2: 89–90
- Simons, D. J., and C. F. Chabris. 1999. Gorillas in our midst: Sustained inattention blindness for dynamic events. *Perception* 28 (9): 1059–1074.
- Smith, L. S., C. Hesp, A. Lutz, J. Mattout, K. Friston, and M. Ramstead. 2020. Towards a formal neurophenomenology of metacognition: modelling meta-awareness, mental action, and attentional control with deep active inference.
- Sterling, P. 2012. Allostasis: a model of predictive regulation. *Physiology & behavior* 106 (1): 5–15.
- Ullman, T. D., E. Spelke, P. Battaglia, and J. B. Tenenbaum. 2017. Mind games: Game engines as an architecture for intuitive physics. *Trends in cognitive sciences* 21 (9): 649–665.
- Van de Cruys, S. 2017. *Affective value in the predictive mind*. 0–0. Frankfurt am Main: MIND Group.
- Ward, D., T. Roberts, and A. Clark. 2011. Knowing what we can do: actions, intentions, and the construction of phenomenal experience. *Synthese* 181 (3): 375–394.
- Whitwell, R. L., A. D. Milner, and M. A. Goodale. 2014. The two visual systems hypothesis: new challenges and insights from visual form agnostic patient DF. *Frontiers in Neurology* 5: 255.
- Wiese, W. 2017. Action is enabled by systematic misrepresentations. *Erkenntnis* 82 (6): 1233–1252.
- Marchi, F., & Hohwy, J. 2020. The intermediate scope of consciousness in the predictive mind. *Erkenntnis* 87, 1–22.
- Bridgeman, B., Kirsh, M. and Sperling, A., 1981. Segregation of cognitive and motor aspects of visual function using induced motion. *Perception & Psychophysics*, 29(4), pp.336-342.
- Goodale, M.A., Pelisson, D. and Prablanc, C., 1986. Large adjustments in visually guided reaching do not depend on vision of the hand or perception of target displacement. *Nature*, 320(6064), pp.748-750.
- O'regan, J.K. and Noë, A., 2001. A sensorimotor account of vision and visual consciousness. *Behavioral and brain sciences*, 24(5), pp.939-973.
- Hesp, C., Smith, R., Parr, T., Allen, M., Friston, K.J. and Ramstead, M.J., 2021. Deeply felt affect: the emergence of valence in deep active inference. *Neural computation*, 33(2), pp.398-446.
- Friston, K. J. 2019. Beyond the desert landscape. In M. Colobmo, E. Irvine, & M. Stapleton (Eds.). *Andy Clark and his critics*. Oxford: Oxford University Press.
- Jackendoff, R. (1987). *Consciousness and the computational mind*. The MIT Press.
- Prinz, J., 2000. A neurofunctional theory of visual consciousness. *Consciousness and Cognition*, 9(2), pp.243-259.
- Koch, C. (2004). *The quest for consciousness*. Engineering and Science, 67(2), 28-34.
- Prinz, J. (2012). *The conscious brain*. Oxford University Press.