



Is free-energy minimisation the mark of the cognitive?

Julian Kiverstein^{1,2}  · Matt Sims³

Received: 24 October 2020 / Accepted: 8 February 2021 / Published online: 1 April 2021
© The Author(s) 2021

Abstract

A mark of the cognitive should allow us to specify theoretical principles for demarcating cognitive from non-cognitive causes of behaviour in organisms. Specific criteria are required to settle the question of when in the evolution of life cognition first emerged. An answer to this question should however avoid two pitfalls. It should avoid overintellectualising the minds of other organisms, ascribing to them cognitive capacities for which they have no need given the lives they lead within the niches they inhabit. But equally it should do justice to the remarkable flexibility and adaptiveness that can be observed in the behaviour of microorganisms that do not have a nervous system. We should resist seeking non-cognitive explanations of behaviour simply because an organism fails to exhibit human-like feats of thinking, reasoning and problem-solving. We will show how Karl Friston's Free-Energy Principle (FEP) can serve as the basis for a mark of the cognitive that avoids the twin pitfalls of overintellectualising or underestimating the cognitive achievements of evolutionarily primitive organisms. The FEP purports to describe principles of organisation that any organism must instantiate if it is to remain well-adapted to its environment. Living systems from plants and microorganisms all the way up to humans act in ways that tend in the long run to minimise free energy. If the FEP provides a mark of the cognitive, as we will argue it does, it mandates that cognition should indeed be ascribed to plants, microorganisms and other organisms that lack a nervous system.

Keywords Free energy principle · Allostatic control · Mark of the cognitive · Minimal cognition · Sensorimotor coordination

✉ Julian Kiverstein
j.d.kiverstein@amsterdamumc.nl

¹ Amsterdam University Medical Centre, Meibergdreef 9, 1105AZ Amsterdam, The Netherlands

² Amsterdam Brain and Cognition, University of Amsterdam, Amsterdam, Netherlands

³ School of Philosophy, Psychology and Language Sciences, University of Edinburgh, George Square, Edinburgh, Scotland

Introduction

Morgan's canon states that comparative psychologists should not "interpret an action as the outcome of the exercise of a higher psychical faculty, if it can be interpreted as the outcome of the exercise of one which stands lower in the psychological scale" (Morgan 1894: p.53). Morgan's advice has been taken to heart by comparative psychologists and rightly so. It is only right that psychologists studying the minds of non-human animals avoid the twin threats of anthropomorphism and intellectualism, ascribing to animals capacities for thinking and reasoning that they have no need for given the environments they inhabit (Barrett 2011). But there is also an opposite danger that arises once psychologists give too much weight to Morgan's canon. This is the risk of oversimplifying the causes of animal behaviour, looking for a non-cognitive explanation when a cognitive explanation is in fact warranted. The risk is that of neglecting the flexibility, adaptivity and complexity of animal behaviour in the search for simple, non-cognitive explanations. Plants and microorganisms exhibit a staggering degree of flexible adaptive agency despite their lack of a nervous system (Müller and di Primio 2001; Hellingwerf 2005; Van Duin et al. 2006; Shapiro 2007; Wolfe et al. 2008; Saigusa et al 2008; Dussutour 2010; Baluška and Mancuso 2013; Wang et al. 2015; Lyon 2015; Gagliano et al. 2016; Novoplansky 2016; Pinto and Mascher 2016; Trewavas 2014; Fulda 2017; Calvo & Friston 2017, Fultot et al. 2019; Salek et al. 2019). These allegedly primitive and "simple" organisms are not "mere automata", they exhibit purposive agency (c.f. Burge 2009; Walsh 2015; Di Paolo, Barandiaran and Buhrmann 2017; Fulda 2017).

Any account of organismic behaviour must avoid the twins pitfalls of either underestimating or overintellectualising their cognitive achievements. To meet this challenge, what we will call a mark of the cognitive is required—theoretical principles that allow one to distinguish cognitive from non-cognitive causes of behaviour.¹ We should resist assuming that we already know roughly what cognition is and can therefore rely on our rough intuitions in making judgements about whether or not an organism's behaviour is caused by cognitive processes. Our rough intuitions about what cognition is are formed in large part on the basis of what has typically been investigated under this heading in cognitive science—namely, human thinking, reasoning and problem-solving. If we hold other organisms to such a standard of possessing human-like thoughts they will inevitably fall short. Alternatively we

¹ One might suggest, in agreement with Andy Clark, that cognitive science has no need for a mark of the cognitive. Clark suggests we can rely on our rough folk-psychological intuitions that tell us when behavioural patterns are indicative of cognition (Clark 2011: p. 451). More recently he has expressed scepticism about the very concept of the cognitive suggesting that the terms "mind" and "cognition" are "confusing and unstable terms that should play no role in mature cognitive science (Clark 2019: p. 294). For the time being, until cognitive science matures, perhaps we can simply make do with calling "cognitive" whatever is studied by cognitive science (Allen 2017). But such an approach would not help us with the challenge we have just described of when to treat the behaviour of organisms as cognitive. Here we are in agreement with Adams (2019) when he suggests it does not help to say "cognition is as cognition does" because one still wants to know which behavioural effects are effects of cognition and which are not.

will make the mistake of overintellectualising the minds of other organisms, flouting Morgan's canon.

In what follows we will use Karl Friston's Free Energy Principle (FEP) to identify a mark of the cognitive that avoids both of these pitfalls (Friston and Stephan 2007; Friston 2010). The FEP purports to describe principles of organisation that any organism must instantiate if it is to remain well-adapted to its environment. It tells us that if an organism is to remain well-adapted to a dynamically changing environment, it must act to minimise an information-theoretic quantity known as free-energy. Living systems from plants and microorganisms all the way up to humans act in ways that tend in the long run to minimise free energy (Friston 2013; Auletta 2013; Calvo and Friston 2017; Ramstead et al. 2018). We will argue for two claims in what follows: first we will show how the FEP can serve as the basis for a mark of the cognitive, demarcating cognitive from non-cognitive processes. Second we will show on this basis that it does indeed make sense to characterise plants, microorganisms and other organisms that lack a nervous system as instantiating cognitive processes.

Our argument unfolds over four sections. We begin in Sect. 1 by introducing the free energy principle (FEP). We interpret the FEP as describing the organising principles of all complex adaptive systems that act to ensure their own continued viability and flourishing in their niche. The FEP as a mathematical description of principles of organisation of complex adaptive systems should be contrasted with a process theory such as predictive processing or active inference which purport to explain how FEP could be biologically realised.² In Sect. 2 we show how Hohwy treats processes of prediction error minimisation on the model of processes of thinking human scientists engage in when formulating true hypotheses on the basis of abductive inference. Hohwy takes processes of prediction-error minimisation to reach all the way down the tree of life to the very earliest appearances of life in single-celled organisms. Understood in this way FEP would suggest a mark of the cognitive as abductive inference. In Sect. 3 we discuss how Hohwy resists drawing such a consequence, reserving the term cognition for organisms that are capable of counterfactual predictive processing (Corcoran et al. 2020). We disagree with Corcoran and colleagues arguing for an understanding of cognition in terms of allostatic control, a position they reject. We show how allostasis is at the core of sensorimotor coordination allowing for the regulation of sensorimotor behaviour we call "allostatic control". In Sect. 4 we provide a variety of examples of how allostatic control is ubiquitous, occurring in primitive organisms that develop a model of their environments. Allostatic control can be described as the process of selecting actions that minimise expected future free energy. It is found everywhere in nature. Thus we

² Hohwy (2020) contrasts the FEP's epistemic status as a principle with that of law of nature that require empirical evidence. Hohwy suggests that it is process theories informed and guided by FEP that rest on empirical evidence, not FEP itself. The reasons for believing FEP follow from the "particular concept of existence of self-organising (nonequilibrium steady state) organisms, together with the mathematics of variational calculus" (ms, p. 9). See also Colombo and Wright (2018) for a good discussion of the status of FEP as what they call a "first-principle".

will conclude on this basis that FEP does provide a theoretically principled basis for ascribing cognitive processes to organisms that do not have a nervous system.³

The free-energy principle: a brief overview

The Free Energy Principle (FEP) states that all self-organising systems that are able to persist over time and thus resist a tendency to increasing disorder will do so by minimising an information-theoretic quantity known as free energy. We will call systems that are able to avoid dispersal and disintegration in their exchanges with the environment “complex adaptive systems.” Biological systems are special instances of complex adaptive systems.⁴ In part by means of a continuous energetic exchange with their external environments, such systems are able to produce and maintain the processes that make them up under conditions of continuous material change. The result of this active self-production is an organism that can persist as an individual through its exchanges of matter and energy with the environment. A cloud of smoke when perturbed by the wind will deform and disperse according to the distribution, magnitude and direction of the mechanical forces acting upon it. Complex adaptive systems by contrast are able to keep key variables within homeostatic bounds when perturbed by their environments. When exposed to high temperatures that are incompatible with their metabolic functioning for example, they will self-regulate (e.g., perspiring, panting, or relocating somewhere cooler) so as to enable their body temperature to remain within a range of values consistent with their functional and structural integrity.

The FEP is a principle in the following sense: it describes mathematically the organisation a complex adaptive system must have if it is to remain in a far from equilibrium steady state and thus continue to exist over time in its interactions with the environment. It is a principle then in the sense that from the mathematics of the FEP one can deduce organisational properties a self-organising system must exhibit if it is to exist. The scope of the FEP is maximally wide, ranging from microscale explanations of the adaptive behaviour of bacteria (Auletta 2013; Friston 2013;

³ Thus we take the argument of our paper to align with proponents of what is called minimal biological cognition as a research programme within cognitive science (see e.g. Van Duin, Keijzer and Franken (2006), Lyon (2006) and Keijzer (2020) and Adams (2018) for sceptical arguments).

⁴ It should perhaps be noted that there are non-living systems that are also correctly described as far from equilibrium steady-state systems. Candle flames, tornadoes and laser beams are classic examples. These systems are dissipative structures whose macroscopic patterns of order spontaneously emerge and are stabilised through ongoing exchanges of matter and energy with the environment. One important difference between organisms and these examples of far from equilibrium open systems is that organisms maintain themselves in a dynamically stable steady state through regulating their interactions with the environment. The examples of non-living systems are not able to actively anticipate and accommodate external perturbations as organisms do. Organisms are able to remain stable through change precisely because they act to anticipate and accommodate sensory perturbations from the outside. Blood cells and fat cells also have some degree of autonomy and may engage in anticipatory regulatory control. However, due to the limited scope of this paper, our analysis will specifically target complex adaptive systems at the scale of the organism. For an account of how microorganisms can be dynamically coupled as constituents of a larger free-energy minimizing unit see (Sims 2020).

Tschantz et al 2020) to the macroscale explanations of variational niche construction (Bruineberg et. al., 2018), natural selection (Campbell 2016) and everything in between (Friston 2011, 2012; Ramstead et al. 2018; Kirchhoff et al. 2018; Ramstead et al. 2020a; for critical discussion see Colombo and Wright 2018). The FEP states that the actions any complex adaptive system should select are those that minimise surprising or improbable sensory exchanges with the environment. Complex adaptive systems avoid global thermodynamic equilibrium—a state in which there is no flow of matter or energy either within the system or between the system and its environment. According to the FEP, they do so by keeping the information theoretic quantity of surprise to a minimum.

Surprise quantifies how unexpected (i.e., improbable) sensory states are given a model of the organism's eco-niche (Tribus 1961).⁵ Our talk of sensory states is short-hand for the states of perceptual systems that are tuned to structure in the eco-niche. We will use the term sensory states in a broad sense to include for instance the states of photoreceptors in plants. Of the range of states a perceptual system can possibly visit, there is a relatively small subset that the system should visit if it is to remain well-tuned to the structure of the organism's eco-niche. We will call the sensory states that belong to this sub-set the organism's "attracting set" (following Friston 2012). As long as the organism remains well-adapted to its environment the sensory states belonging to its attracting set are those it should tend to be attracted towards. The existence of this attracting set allows us to construct a probability distribution for the organism's sensory states which represents the probability of finding an organism in a given sensory state. This distribution can then allow observers to measure what we will call the sensory state's "surprise value"—the probability of finding the perceptual system in a sensory state when sampling the system at any random time. Very roughly, we can say a sensory state s has a low surprise value when the probability of an organism regularly revisiting s over time is high. While s will have a high surprise value when the probability of an organism finding itself in s over the course of its lifetime is low. If the organism maintains a distribution over its attracting set with low surprisal, it will remain well-adapted to its eco-niche, and will avoid sensory states that are a threat to its integrity, or that lead it to be maladapted in some other way.

Surprise, however, is a computationally intractable quantity for the organism. It may perhaps be knowable to an observer of an organism but it is not a quantity the organism itself can evaluate. Doing so would require the organism to have access to and make inferences over all of the possible states that it could sample. According to FEP, the organism's way around this problem is to use free energy (FE) as a stand in for surprise value. FE is a function of the organism's internal dynamics and its sensory and active states—the states of its perception and action systems. FE is also an upper bound on (i.e. it is always greater than or equal to)

⁵ More formally, surprise refers to the negative log probability of observing sensory and active states conditioned on a model or equivalently the negative log marginal likelihood. In Bayesian statistics the marginal likelihood is known as Bayesian model evidence. Surprisal is formally expressed as: $-\ln P(s, a | m)$. Here s refers to sensory states, a to action states and m to the model. Conversely, Bayesian model evidence is: $\ln P(s, a | m)$.

surprise value and can thus approximate this latter quantity (Friston 2010; Friston and Stephan 2007). Hence by regulating its internal states and the states of its perception and action systems, the organism can perform gradient descent on FE, and is thereby able to implicitly minimize surprise. The organism can regulate its internal states by changing them in response to external perturbations so as to maintain its organisation. It can also regulate its sensory and active states, moving to bring about the sensory states it expects given the model it entails. An organism that acts in this way to minimise free energy will thereby also avoid sensory states with a high surprise-value that are a threat to its remaining well-adapted to its ecoiniche.

In the literature on FEP, FE is standardly taken to be a “function of probabilistic beliefs encoded by an organism’s internal (i.e. neural) states” (Ramstead et al 2018: p.4). The internal states of an organism are thus taken to parameterise probabilistic densities over states of affairs in the external environment. We should interpret the notion of belief in play in FEP in inferential and probabilistic terms. When we talked above about the organism changing its internal dynamics in response to external perturbations this can be interpreted as a process that approximates Bayesian inference. Bayes’ rule tells us how to rationally change our beliefs when they come into conflict with evidence. It tells us to update the probability of a given belief by considering the product of the likelihood (the probability of the evidence given this belief) and prior probability of the belief (which should sum to 1). Bayes’ rule tells us to infer the belief that has the highest posterior probability. Bayesian inference is computationally complex and often intractable in real world examples of belief update (Wiese and Metzinger 2017). There are simply too many possible beliefs to be evaluated for a belief update system to be able to calculate which of these possible beliefs has the highest posterior probability. FEP describes a computationally tractable method for solving this problem. So long as a system changes its internal dynamics in response to external perturbations in such a way as to minimise FE it will do a good job of approximating belief update in accordance with Bayes’ rule (Cf. Hohwy 2020).

Still this leaves us with the question of the relation between the organism’s internal dynamics and the probabilistic beliefs it is ascribed under FEP. We will assess the following two options for addressing this question:

1. Adopt an instrumentalist reading of the concept of prior probabilistic belief as it figures in the FEP.
2. Offer a realist account of probabilistic belief states as this concept applies to organisms in general from prokaryotes to mammals.

The sensory states that form an attracting set for an organism, and relative to which free energy is computed, define a generative model. We have seen above how when an agent gathers new sensory evidence it must combine a likelihood function (a probabilistic mapping from hidden states of the world and their dynamics x to sensory inputs y) with its prior beliefs (a probability distribution that predicts possible states of the world over time x). These two probability

distributions (the prior beliefs and the likelihood function) are referred to as the “generative model”. The likelihood and priors are described as a *generative* model because they can be interpreted as mapping how sensory inputs y are believed to be generated by states x of the environment. Given some sensory observations the generative model is used to compute the posterior probability of a possible state of the world that is the cause of those observations. Should the generative model in relation to which free energy is defined be understood instrumentally or realistically?

Instrumentalism about the generative model in FEP (and about Bayesian theories of cognition more generally) is emerging as a popular option in the recent literature (see e.g. Colombo & Series 2012; Colombo et al. 2018; Van Es 2020; Ramstead et al. 2020a, b; Baltieri et al. 2020). Van Es (2020) argues for instance that organisms may well behave *as if* they approximate the norms of Bayesian probabilistic inference but they need not *really* engage in any inference, and nor need they *really* encode probability densities in their internal states. He makes his argument by distinguishing between a model used by a scientist to explain and predict the behaviour of a system of interest, and a model as physically implemented by the internal dynamics of the system under investigation. He suggests FEP is best understood as a modelling framework used by scientists to explain the self-organising dynamics of complex adaptive systems that persist over time (Cf. Andrews 2020; Bruineberg et al. 2020).

We follow Friston in conceiving of organisms as becoming models of the structure of their eco-niches over the course of their learning histories (Friston 2012: p. 2101). The generative model is not something the organism encodes in its nervous system (if it has one). As Bruineberg et al. (2018) have argued the generative model is “entailed” by the organism’s internal dynamics that form in its selective coupling to its niche (cf. Friston 2011, 2012). The conditional dependencies the generative model maps—the joint probability distribution over sensory states, actions and interacting environmental causes—are a consequence of the actions the organism regularly undertakes in its niche (c.f. Ramstead et al. 2019: p. 18; Bruineberg et al. 2016). These dependencies are instantiated in the organism’s internal dynamics as they self-organise in its active coupling with its niche. By “internal dynamics” we mean the relatively stable biological organisation an organism generates and maintains through a history of practical engagement with the environment. The organism is a “model” of the structure of its niche in the sense that it instantiates a biological organisation that complements the structure of its niche. Thus we are committed at minimum to the claim that the internal dynamics an organism develops stand in a relation of entailment to a generative model of its environment. There is however an open question remaining about how to understand this relation of entailment, and what more specifically is entailed by the organism’s internal dynamics.

This leads us to a second question of how to interpret the notion of probabilistic prior beliefs that FEP would seem to imply organisms should somehow realise in their internal dynamics. We will give an account of the generative model below that is consistent with instrumentalism about belief-talk. Does it follow that a defender of the FEP should also be instrumentalist about the generative model? We do not see how a more general instrumentalism about the generative model could be consistent

with the characterisation we have given of FEP in this section. We have claimed that FEP purports to deduce the organising principles common to complex adaptive systems. The FEP, on this reading, does not claim that it is useful for explanatory purposes to model complex adaptive systems as if they satisfied these organising principles. It purports to tell us how complex adaptive systems would need to be organised to remain well-adapted to their environments and persist over time. We conclude then that a commitment to the existence of generative models is non-negotiable for proponents of FEP. Should we then also be realists in taking the internal states of organisms to parameterise probabilistic beliefs about external states of the environment? Recall that the generative model has this name because it is posited to capture the organism's beliefs about the causes of its sensory states. We take up this question of the status of these beliefs in the next section.

Realism about Bayesian organisms

Hohwy defends a realist position about belief talk in the context of FEP (Hohwy 2013, 2015, 2019; Kiefer and Hohwy 2017). He has argued that any system that can minimise long-term prediction error will also approximate Bayesian inference. It will operate with prior beliefs based in part on its past learning. When it accommodates its sensory input in such a way as to minimise prediction errors this can be thought of in terms of inference to the best explanation. What the subpersonal processes that make up the organism unconsciously infer is a set of hypotheses that best explain its sensory inputs and their distal hidden causes. What makes a set of hypotheses the best explanation of the available evidence is that the hypothesis exemplifies precision, simplicity and accuracy. As we saw above Hohwy takes it to be plausible that the same process of prediction error minimisation through inference to the best explanation that is to be found in the human brain may also be found in evolutionary primitive organisms such as *E. Coli* (Hohwy 2019: p. 201). It follows that every complex adaptive system must build a model of its environment, perhaps in part through learning, and put this model to work to do abductive inference. What the model targets are the hidden environmental and bodily causes of the sensory states the organism can expect to repeatedly revisit over time given its phenotype. On the basis of this model, they selectively sample sensory states in order to fulfill these expectations. According to Hohwy, the difference in complexity between plants and bacteria and other organisms that embody a nervous system is that the model they construct over time lacks a deep temporal and causal structure. (We return to this point in Sect. 3.) But all lifeforms, as free energy minimising systems, must do inference to the best explanation.

In Hohwy's view, free energy minimisation is biologically realised by processes of unconscious abductive inference. Any living system owes its continued existence to processes that approximate inference to the best explanation. This claim seems to us to flout Morgan's canon. Hohwy seems to invoke sophisticated human-like capacities to explain the behaviour of non-human organisms. Scientists use inference to the best explanation to select among competing hypotheses the one that is the most fruitful, simple and precise in terms of the data it can explain. If Hohwy is right

bacteria, albeit on a much smaller scale, are also capable of performing sub-personal and unconscious inferences that are approximations of the personal-level inferences made by scientists. But this is to overintellectualise processes of free energy minimisations. Other organisms (and indeed humans when they are not doing science) have no need for constructing theories about the world they live in. They need to figure out how to act in ways that are good enough most of the time for them to get by given their ways of life within a given ecological niche. But there are familiar arguments for distinguishing action guiding processes that do a good enough job most of the time from belief-forming processes that aim at the truth (Varela et al. 1991, ch.9; Stich 1990; Burge 2010; Hutto and Myin 2017). Second, an organism can accomplish the adaptive control of their actions without constructing theories about the causal structure of their worlds. We say more on this second point in Sect. 3 below.

Andy Clark has argued for a more restrictive usage of the concept of generative model. He reserves what he calls “predictive processing” for organisms that instantiate a cognitive architecture that is characterised by hierarchical message-passing and top-down prediction. He argues that humans undoubtedly rely upon rich internal models, albeit models that work in the service of efficient, low-cost, high-gain, action selection and guidance. Other animals can make do without such rich models relying only on quick and dirty processing or “genetically preconfigured responses to sensory perturbations” (Clark 2017: p. 6). Thus, Clark proposes to make a distinction between advanced cognisers that encode in their brains explicit expectations that are used to generate top-down predictions and those organisms that rely in some sense on implicit creature-specific expectations in perception and action.

For Clark, it is only those creatures that are capable of learning neurally encoded explicit expectations that *really* instantiate a generative model in their internal dynamics. Friston’s (and our) talk of organisms being a model of their environment is best construed instrumentally as Van Es (2020) and others have suggested. It may well be useful for modelling purposes to treat the fish as if it embodied in its morphology expectations about the hydrodynamics of its aquatic environment. However, if Clark is right such talk is only warranted insofar as it helps us to predict and explain the fish’s behaviour. Strictly speaking, to have a generative model calls for neurally encoded expectations that can be used to generate top-down prediction in a hierarchically organised cognitive architecture.

Clark would thus reject the claim that Hohwy makes that because bacteria and other organisms that lack a nervous system are free energy minimisers, these life-forms must be capable of abductive inference. He can thereby avoid Hohwy’s over-intellectualisation of the cognitive capacities of such organisms. That being said, Clark has very little to say about the cognitive powers of organisms that fall on the non-predictive processing side of his divide. Presumably he thinks these organisms fallback on some form of “genetically preconfigured response” (Clark 2017). Do such responses also count as cognitive? Clark doesn’t say, and would if pressed no doubt express scepticism about whether the question really matters (Clark 2011, 2019, see footnote 1). We disagree. There is a genuinely important question to be settled here about where cognition gets off the ground in the evolution of life. Recall in our introduction how we pointed to the real risk of oversimplifying the impressive cognitive achievements of evolutionarily primitive organisms in the search for

non-cognitive explanations of their behaviour.⁶ We will argue next that the Free Energy Principle (FEP) underwrites a mark of cognition that can strike the right balance, neither overintellectualising nor underestimating the cognitive capacities of other non-human lifeforms.

Where does cognition begin in life?

In their co-authored paper published in this special issue, Corcoran and colleagues address exactly the question we are now grappling with of how to draw the line between the cognitive and non-cognitive (Corcoran, Pezzulo and Hohwy 2020). They suggest that a capacity for decoupled representation is required for cognition: “only those biological systems capable of engaging in fully detached modes of representation, and of exploiting such representations for the purposes of uncertainty reduction, count as cognitive agents” (*Op Cit.*). They thereby propose to restrict cognition to creatures that are capable of what they call counterfactual active inference.

Corcoran and colleagues consider distinguishing “relatively primitive creatures (like *E. coli* and other unicellular organisms) creatures with some degree of hierarchical depth (like reptiles and fish), and animals that demonstrate evidence of counterfactual sensitivity” such as rodents, corvids and primates (p. 39). The notion of hierarchical depth they appeal to relates to the generative model with the distinct layers of processing regulating action under different conditions of uncertainty and over increasingly larger spatial and temporal scales of sophistication. As the hierarchical depth of the model increases, so the decision-making capacity of the creature is increasingly decoupled from the creature’s immediate sensory input. We could thus extrapolate from Corcoran and colleagues an account of the evolution of a capacity for decoupled representation in terms of the hierarchical depth of the generative model an organism embodies.

Primitive organisms instantiate a model that is used to infer the presence of nutrients or noxious stimuli. These creatures operate with models that are shallow in their spatial and temporal depth. The generative model of a primitive creature cannot be used to build up a map of conditions encountered in the environment previously in foraging so as to find their way to possible sources of nutrition again in the future. Nor can the model be used to anticipate the sensory consequences of an action policy because these creatures are unable to learn about what structures their sensory states in their niche. All the creature can do is engage in what is called “closed-loop control” moving more or less randomly when a prediction error (i.e. a divergence from its expected sensory state) occurs (*Op cit.* p.25).⁷ Fish and reptiles can do more

⁶ This is not to claim that such lifeforms rely upon rich, reconstructive, truth-seeking models of their ecological niches. Our point is simply that we should take care not to underestimate the cognitive achievements of organisms that are able to make do without such reconstructive models of their environments.

⁷ We will see later how this is something of a caricature of what *E. coli* can do. To be fair to Corcoran and colleagues they are explicit in saying that they mean to describe a highly simplified model of primitive organisms here and they allow that the behaviour of bacteria could very well turn out to be much more complex. Later in their paper they suggest that “all extant lifeforms instantiate some form of

with the models they develop. They can take into account the likely consequences of actions in the future based on their past experience and select actions that have the lowest risk (i.e. expected uncertainty). They can make use of knowledge about how their sensory states change over longer intervals of time and take proactive measures to keep their sensory states within a range of preferred values. For example, they may be able to budget the body's energy needs so more energy is available during the day for foraging than at night before sleep. Corcoran and colleagues suggest that rodents, corvids and primates make use of models with the deepest temporal reach. They are capable of some form of counterfactual active inference. They can simulate the sensory states the organism would sample were it to embark on a sequence of actions under a variety of possible but non-actual environmental conditions. They can, in this way, manage volatility or unexpected uncertainty allowing them to differentiate risk that is due to an action policy from risk that is due to change in environmental dynamics.

We agree with Corcoran and colleagues that it is important not to conflate life with cognition. We should resist drawing any simple relation of equivalence between life and cognition. Positing such a relation of equivalence would not help us to understand the relationship between life and cognition (Van Duijn et al 2006: p. 160). We should instead look for gradations in complexity of cognition in different lifeforms, a “shading-off” of cognition into more basic biological processes (Godfrey-Smith 2002, 2016). However we disagree with the sharp discontinuity Corcoran and colleagues insist upon between “proto-cognition”, which they concede may well be present in all living systems, and “genuine cognition”, which they take to be characterised by counterfactual inference. Such a sharp discontinuity implies that organisms that can learn, make decisions and form memories are not, just by dint of having these capacities, cognisers. They need in addition to implement a generative model of sufficient sophistication to allow them to engage in counterfactual inference to count as cognising. This strikes us as again an over-intellectualisation of cognition. While not unprincipled, we take it to be nevertheless unwarranted, and certainly not implied by the FEP.

On the contrary we take the FEP to provide strong support for, even if it does not imply, what Lyon (2006) has called a “biogenic view” of cognition.⁸ The biogenic view starts from principles of biological organisation to provide a principle basis for defining the border between the cognitive and the non-cognitive. Lyon used the word “genic” to describe this view because she is concerned with the origins or beginnings of cognition in the evolution of life. Searching for the origins of cognition means starting one's search with the simplest of organisms in terms of

Footnote 7 (continued)

allostatic architecture” (p. 39). We agree with them on this point and offer some worked out examples to illustrate and support this point in the final section of our paper below.

⁸ Corcoran and colleagues deny that there is any implication from FEP to strong life-mind continuity—the view that wherever there is life there is also mind (Kirchhoff and Froese 2017). They defend a weak continuity view according to which anything that has a mind will also be alive. We will argue by contrast that the organisational principles of life as characterised by the FEP provide a mark of cognition—criteria for demarcating the cognitive from the non-cognitive without simply conflating life and cognition.

their organisation (but not their behaviour as we will see later)—prokaryotes. On the biogenic view cognition is a natural biological function that allows organisms to flourish and make a life in their particular econiches. We should not assume without argument that cognition requires a nervous system. Once we look for the origins of cognition we should not rule out the possibility that organisms that lack a nervous system such as plants, fungi, archea and bacteria can nevertheless cognise.

We take as our starting point Van Duijn et al. (2006) proposal that cognition got started with the capacity for sensorimotor coordination (Cf. Kiverstein 2020).⁹ Organisms such as plants and bacteria that lack a nervous system are nevertheless capable of sensorimotor coordination with their econiche. Cognition didn't start with metabolic processes but with the coordination of sensorimotor behaviour with the environment. Sensorimotor coordination processes, as Keijzer nicely notes, “constitute major cognitive achievements, something that has become self-evident when one tries to mimic similar capacities in robotic devices (Pfeifer and Bonggaard 2007)” (Keijzer 2020: p.6). Plants, fungi, archea and bacteria do not possess brains but they all exhibit behaviours that are well-adapted to their circumstances. Slime molds for instance have been shown to be capable of navigating mazes (Ried and Latty 2016), a cognitive achievement that requires learning and memory. Plant growth involves decision-making that requires balancing different constraints and that distinguish edible from non-edible, or preferred from less-preferred sources of nutrition (Trewavas 2014; Calvo and Keijzer 2011; Linson and Calvo 2020). These organisms engage in a rich range of behaviours that coordinate with the constraints of their circumstances and that transcend metabolic-related functions.

But if cognition got off the ground with sensorimotor coordination this gives us alternative conceptual tools for thinking about the question we posed earlier (at the end of Sect. 1) of how to think about the probabilistic expectations relative to which FE is defined. The realist view as defended by Hohwy takes internal dynamics to parameterise probabilistic beliefs that represent hidden external causes of sensory input. On this realist view, the brain operates according to the norms of Bayesian probabilistic inference, aiming to represent truthfully the causal structure of a mind-independent external reality. We are suggesting by contrast that internal dynamics parameterise expectations about the sensory consequences of action. Instead of conceiving of these expectations in terms of belief, as the realist proposes, we propose to think of expectations in terms of what Susan Hurley called sensorimotor dynamics (Hurley 2010; Cf. Noë 2004). The parameters of the organism's internal dynamics are the parameters of the organism's perception and action systems that are tuned to the structure of its econiche [i.e. the action possibilities or affordances the econiche provides in part because of the organism's own actions (Bruineberg et al. 2016)]. In acting according to the imperative to minimise free energy the organism does not aim at true posterior beliefs but at tuning its action systems to the structure of its econiche in ways that support adaptive behaviour.

⁹ In this paper we will however widen van Duijn, Keijzer and Franken's conception sensorimotor coordination to include not only behaviour that is movement based (e.g., forms of taxis, locomotion, etc.) but also to include integrated biochemical behaviour that can be ascribed to the entire organism.

According to the good regulator theorem of Conant and Ashby (1970) the internal dynamics an organism embodies must be a good regulator of the environment of which it is a part. The generative model should therefore be interpreted as a control system. Crucially control takes place without any inner executive that is coordinating, and micro-managing the control processes. Control is self-organising. The more complex the ecosystem of an organism and the affordances it makes available, the more complex the control system the organism will need to remain well-adapted to this niche (cf. Cisek 1999; Anderson 2017; Seth 2015, 2019). A good regulator is one that finds the right balance between maximising accuracy and minimising complexity. Accuracy here should be understood not in terms of the norm of aiming at the truth but in terms of anticipating the sensory consequences of acting in such a way as to remain well-adapted to a dynamic environment. Complexity we take to be ecological balance: the complementarity or alignment between the perceptual and motor skills of the organism and the affordances of its econiche. Minimising complexity here is achieved when the organism finds the right relation of mutual fit between its internal dynamics and the external dynamics of the environment.

The kind of control the generative model underwrites is what we will call “allostatic control”. We have seen how complex adaptive systems expect to occupy sensory states belonging to their attracting set and act to fulfill this expectation. The attracting set will include global physiological, visceral and vascular conditions of the organism’s body such as cardiac signals, states of the gut and viscera, glucose and oxygen levels. Allostatic control is anticipatory, as we will see in some detail in the next section. The organism anticipates deviations from sensory states belonging to the attracting set before they arise and takes appropriate precautionary actions to maintain itself in a non-equilibrium steady-state. We suggest that allostatic control *drives* sensorimotor behaviour ensuring the organism remains well-adapted to its environment. Allostatic control is a form of affect-based regulation of behaviour.¹⁰ We can think of the maintenance of the organism in a far from equilibrium steady-state as a fundamental concern—a goal or purpose for which the organism strives. Insofar as the organism has this goal, anticipated deviations from expected physiological conditions are registered by the organism as affective states. The organism can regulate its sensorimotor engagement with the environment by using its affective states to initiate goal-directed, purposeful movement. Minimally, this will take the form of approach-avoidance behaviours, either moving away from what is potentially harmful or towards what would improve the organism’s material conditions in its environment. Thus, sensorimotor coordination can be seen as a consequence, at least in part, of affective processes of allostatic control.

We suggest that this biological strategy may be found in a wide range of living systems; it is not only neuronal organisms that make use of allostatic control

¹⁰ When we describe these anticipated deviations as registering with the organism affectively we do not mean to imply that such deviations are consciously felt, implying some form of sentience in all organisms capable of allostatic control. We take allostatic control to be affective in the sense that anticipated deviations from homeostatic setpoints are a matter of concern to the organism, which in our view justifies describing allostatic control as a form of affect-based regulation (Cf. Lyon 2015). We use affect here in a sense that is neutral on whether the affective states are conscious. The question of when consciousness first made an appearance in life is beyond the scope of this paper.

to achieve sensorimotor coordination with the environment, but also basal organisms (Cf. Alfieri 2008). It is an implication of the kind of self-organizing dynamics falling out of FEP that all organisms will possess capacities for allostatic control.¹¹ Consider in this light Calvo and Friston's example of a plant that grows along a seashore (Calvo and Friston 2017: p. 4). The presence of salt in the soil is a part of the econiche of this plant. But too much salt is associated with free energy for the plant—it is not what the plant expects to encounter. One way the plant might respond is by learning to expect salty encounters and adjusting its morphology accordingly (e.g., osmotic adjustment, efflux, hardening, etc.). Part of this learning however will probably be for the roots of the plant to explore, selectively sampling areas of soil where the concentration of salt is low. Exploration itself is a metabolically costly behaviour given that root nutation (i.e., directed and non-directed root growth) requires both time and the use of energetic resources. It stands to reason that plants cannot afford to invest in high risk behaviour that lands them in stress inducing conditions—surprising conditions which they do not expect to occupy. In order to avoid both salty stress and unnecessary use of metabolic resources, “plants must adjust with respect to future conditions if their behaviour is to remain adaptive” (Calvo and Friston 2017, p. 4). That is to say plants should engage in a strategy of looking ahead, prospectively directing their behaviour to avoid anticipated salt concentration increases *and* investing in metabolic energy expenditure according to the anticipated sensory outcomes of such an expenditure. It is by engaging in this kind of prospective strategy that a plant continues to occupy the attracting set of states that define its existence, harvesting evidence for its generative model.

Calvo and Friston's useful example helps to expose an important aspect of plant adaptive behaviour: Although there is *short-term* metabolic cost associated with nutation, such behaviours are nonetheless invested in because their anticipated consequences allow the plant to occupy states that it expects to occupy over the *long run*. Plants are motivated to incur the immediate metabolic cost and possible short-term divergence from other specific homeostatic expectations defined under the plant's generative model. They incur these costs because they are motivated by the outcomes of their behaviour such as avoiding high salt concentrations and obtaining nutrients. These outcomes are what the plant expects given the model of its ecological niche the plant develops and the plant engages in nutation in order to fulfill these expectations.

In the next section we will look more closely at the notion of allostatic control and its relation to homeostasis from the perspective of FEP. We then turn to some

¹¹ An anonymous reviewer asked us if it is an implication of our view that organisms engage in what we are calling allostatic control all of the time. One might think for instance that environmental volatility is dynamic. Organisms need to have recourse to allostatic control only in those periods in which the environment conditions are predicted to be threatening to an organism's homeostasis. When environmental conditions are predicted to present no challenge to long-term homeostasis, an organism's adaptation to its environment may occur merely via the engagement of reflex-like (non-anticipatory) responses. We've argued above that allostatic control occurs as part and parcel of sensorimotor coordination. We would see the engagement of reflex-like responses that maintain homeostasis and what we are calling allostatic control as part of the process of sensorimotor coordination, though they may unfold at distinct layers of the generative model (as is the case in Stephan and colleagues model of allostasis—see Stephan et al. 2016).

empirical studies which buttress the claim that at least some of the behaviour of non-neural organisms is driven by allostatic control. We will look at various instances of allostatic behavioural control in microorganisms and plants. Our aim in doing so is to support the idea that varying degrees of allostatic control may be a ubiquitous feature of living systems and thus may serve as the basis for ascribing cognition to all complex adaptive systems.

Allostatic control: the mark of the cognitive?

Allostatic control is best understood in relation to the concept of homeostasis as introduced by the physiologist Walter Cannon (1929). Cannon used the notion of homeostasis to describe the maintenance of the conditions of an organism's internal milieu within a constant range of values. Homeostasis has the consequence that the organism maintains balance with its environment, despite the influence of external environmental stressors. Examples of such internal conditions include, for example, blood pressure, heart rate, osmotic pressure, blood glucose and blood oxygenation levels. Remaining alive requires that a system regulates itself so as to ensure that these homeostatic variables remain *around* a constant preferred value (i.e., near a homeostatic setpoint). Think of the fact that the body temperature of humans may vary 0.5 degrees above and below the mean of 37.0 °C without representing a challenge to viability. Homeostatic variables vary in the range of values that are compatible with life. Blood glucose levels may vary to a larger degree over time without having harmful effects.¹² Blood oxygenation levels tolerate very little variation. Given that living systems may encounter wide ranges of environmental conditions and physiological stressors, maintaining homeostasis will often prove to be extremely challenging. Typically, it is assumed this challenge is met through purely *reactive* forms of self-governed regulation.¹³ When a divergence from setpoints is registered, regulatory mechanisms 'kick-in', engaging homeostatic reflex arcs to counteract the disturbance. This restoration of homeostasis is accomplished by subsystems such as the autonomic, immune, endocrine or motor systems. The autonomic system, for example, contributes to the stabilization of respiratory rate and hence blood oxygenation variables, rather than the living system to which that autonomic system belongs.

Allostasis can be defined as "the proactive deployment of behavior, guided by predictions from a model, in order to avoid dyshomeostatic future states" (Stephan et al 2016, p. 28; c.f. Sterling 2012). Allostatic control consists in prospective behaviour directed at avoiding the anticipated divergence from homeostatic setpoints. Homeostasis may successfully minimize *encountered* free-energy only after a homeostatic setpoint has been breached. Allostasis involves the selection of sequences of

¹² To see this, think of how some people may fast for days on end, but they may only do without oxygen intake for minutes at a time.

¹³ To be fair, there are other more recent conceptions of homeostasis which see it as involving learning and anticipation. See Davis (2016). However, in what follows we will continue to use the term homeostasis as it was originally introduced by Cannon.

actions that minimise *expected* free energy (i.e. dyshomeostasis expected as a consequence of action (cf. Friston et al. 2015; Schwartenbeck et al. 2019). Prospective behaviour is crucial for organisms, the continued viability of which remains risky and hostage to good fortune so long as they rely just on passively reacting to environmental changes (Rosen 2012; Lyon, 2006; Goodson 2003). Consider the following analogy: Prior to hiking through a hot desert, it is a good expenditure of energetic resources to locate water and fill water flasks in the desert-bordering village even if one is not currently thirsty. Better to take prospective action now than wait until one's only option is to forage for water in the desert and risk dehydration. Thus, despite its energy costs, the action-policy of loading up on water before entering the desert outweighs the dyshomeostatic effects of severe dehydration that would likely result from failing to act proactively.¹⁴

To the extent that allostatic control requires the selection of action policies that minimise *expected* free energy, one might think this kind of action control could only possibly occur in organisms with nervous systems that can forecast the consequences of their actions. Corcoran and colleagues seem to rely on this type of reasoning. They suggest that to select action policies that minimise risk (i.e. expected uncertainty associated with an action-outcome), the organism must be capable of imagining different courses of action and their outcomes, testing out these actions in the safety of their imagination (ms, p. 36; c.f. Craik 1943; Dennett 1995). The capacity to disengage from the present and imagine different possible futures is however clearly beyond the cognitive capabilities of organisms that do not have a brain.

We have two answers to this challenge. First we will argue that learning an action-oriented model of an environment already requires minimising expected free energy. Second we will offer a number of examples that demonstrate how development of a model of the environment through gene regulation and phenotypic plasticity can also be viewed as minimising expected free energy. Thus, we dispute the claim that dealing with expected uncertainty requires sophisticated capacities for counterfactual reasoning.

With regards to our first claim, consider a recent computational model of chemotaxis (Tschantz et al. 2020). Tschantz and colleagues take up the question of how an organism can learn from experience an action-oriented model of its environment. Organisms actively sample information from their environments based on their pragmatic goals and purposes. They thereby avoid the need for building a fully reconstructive and accurate model of their environments, as argued in the previous section. However a problem remains? How is it then that organisms avoid what the authors call "bad-bootstraps" in which they sample sub-optimal observations which are then used to update the model thereby leading the agent to continue to sample sub-optimal observations? The organism needs to perform not just pragmatic goal-directed actions but also epistemic actions in which it actively forages and seeks out new information. But now the organism will be faced again with the problem of working out how the information it is sampling is relevant to its actions. To learn a

¹⁴ Strictly speaking, thirst does not signify the state of dehydration but occurs prior to dehydration. This allows for one to act before dehydration prevents one from being able to forage for water. A similar relation holds between experiencing hunger pangs and blood glucose deficiency.

model that can function as a good-regulator it seems the organism will need some techniques for balancing what are called “pragmatic” and “epistemic actions”, working out when it is appropriate to exploit the information it has to hand and when to explore for new information. They must in other words have some means of evaluating when a pragmatic or an epistemic action is worth performing. We have suggested in our discussion of allostatic control that the currency that is used to evaluate actions is expected free energy. The authors show that in their model of chemotaxis, the agents that selected between pragmatic goal-directed actions and epistemic actions with the aim of minimising expected free energy were the agents that did the best in learning an action-oriented model of their niche. The agents in their model that selected actions with the aim of minimising expected FE avoided bad bootstraps on the one hand but they also avoided learning accurate and exhaustive models that were no longer useful for purposes of regulating and guiding action.

Epistemic actions are actions in which the agent risks exploring parts of the environment that are currently unknown because the benefits of doing so are gains in information where these gains in information are to be understood in terms of minimising future risk. The agent performs actions that may well lead to free energy in the short term but they do so because they do better at minimising free energy in the long-run. We will argue next there are good reasons to think that non-neural organisms do indeed engage in such behaviours that temporarily take them away from homeostasis with the aim of minimising expected free energy. We take this to show that these organisms are engaging in allostatic control. Given the arguments of this paper, they are exhibiting cognition-based behaviours.

Example 1 Allostatic control in *Escherichia coli* The bacterium *E. coli* offers our first striking example of an instance in which adaptive behaviour takes into account future expected risk. During its life, *E. coli* occupy a multitude of niches, the mammalian gastrointestinal tract being the most frequent (Mitchell et al. 2009). It goes without saying that these niches have extremely different physiochemical properties. When *E. coli* are ingested, moving from outside environments to inside the oral cavity, they encounter both a transition from aerobic to anaerobic environmental conditions and an increase in ambient temperature. Failing to adapt to the oxygen downshift means death for such organisms. Tagkopoulos et al. (2008) emulated the changing temperature and oxygen conditions that *E. coli* encounter when moving from outside to inside the gastrointestinal tract in an artificial environment (i.e., a bioreactor). They hypothesized that if bacterial regulatory behaviour is merely a homeostatic reaction, then bacteria would wait until detecting oxygen downshift to repress aerobic respiration. However, if *E. coli* are capable of anticipatory behaviour, the repression of aerobic respiration might occur proactively, prior to detecting decreased oxygen levels. Remarkably, Tagkopoulos et al. found that *E. coli* use a temperature upshift from 25 to 37 °C—the kind that would typically accompany being ingested—as a cue to repress aerobic respiration.

Tagkopoulos and colleagues posit that this cellular regulatory response is underwritten by a form of learning over evolutionary timescales that is akin to associative learning occurring over ontogenetic timescales. To test this hypothesis, they evolved *E. coli* in conditions that run counter to those found in nature, exposing bacteria to

a temperature upshift followed by an oxygen upshift. They found that *E. coli* when exposed to such counter conditions over time learnt to not repress aerobic respiration, supporting the notion that the anticipatory behavior exhibited by these bacteria is a consequence of their learning the environmental contingencies to which they are exposed. This is to say that because respiration pathway reprogramming is decoupleable from the typical order of environmental cues “the originally observed correlated responses to temperature up-shift and oxygen down-shift cannot be due to hard biochemical constraints, but rather is a reflection of a common response to correlated changes in temperature and oxygen that has evolved over geological time-scales” (Tagkopoulos et al. 2008, p. 5).

We suggest that this behaviour is an example of allostatic control underwritten by their becoming a model of their environment.¹⁵ This is evident from the fact that *E. coli* initiates adaptive regulatory behaviour geared toward future anaerobic conditions *whilst still in the presence of oxygen*, when these bacteria are tuned to the dynamics of an environment in which oxygen downshifts follow temperature upshifts. A bacterium’s finding itself in an anaerobic environment while continuing to expect the presence of oxygen is tantamount to its incurring free energy through its actions. Thus, preadaptive respiration repression here can be thought of as the bacterium minimising expected FE and hence minimizing the risk of finding itself in metabolically poor conditions. The gain from avoiding anaerobic stress exceeds the cost of its anticipatory respiration repression. Taking the diagnostic tool into account, this ability of *E. coli* to invest its current homeostatic resources so as to bring about behaviour that compensates for anticipated future environmental conditions is, we suggest, a form of allostatic control.

Example 2 Allostatic Control in Wild-type Yeast Mitchell et al. (2009) tested the ability of wild yeast (*Saccharomyces cerevisiae*) to anticipate environmental change and compensate for it with prospective behaviour. They investigated the transition in physiochemical conditions during the process of the brewing of alcohol in which these eukaryotes must adapt from yeast-preferred fermentation conditions to subsequent oxidative respiration conditions. During fermentation, due to rise in temperature, *S. cerevisiae* are exposed to heat stress in addition to ethanol stress. Subsequent oxidative respiration generates oxygen radicals which may have deleterious effects upon the cells (Maris et al. 2001). There is thus a causal trajectory from stresses resulting from fermentation to those resulting from respiration. Mitchell et al. were able to demonstrate that *S. cerevisiae* use early heat shock as a cue for encountering subsequent oxidative stress. Modification of regulatory networks due to the detection of heat stress allows yeasts to physiologically adapt to oxidative stress prior to encountering oxidative conditions. Strikingly, heat-induced-genes are at best neutral or at worst maladaptive under heat shock. But heat-induced genes are essential

¹⁵ Tagkopoulos et al. interpret their results similarly. They write: “More generally, the correlation-structure of the environment can be internalized as a probabilistic model in the high-dimensional space of an organism’s complete sensory perception. As such, the very organization of microbial regulatory networks may, in large part, represent the physical instantiation of this probabilistic model. (2008, p. 6).

under conditions of oxidative stress. Mitchell and colleagues were therefore able to confirm “the claim that an early preparation to oxidative stress during heat shock is adaptive in the system” (Mitchell et al. 2009, p. 223). This anticipatory behaviour depends upon both a regular temporal order of unfolding events in the yeasts’ environment and that such order has been captured by the organization of the yeasts’ regulatory networks. In other words, *S. cerevisiae* becomes a model of its environment over the course of its development, which allows the yeast to minimise risk. Such anticipatory physiological regulation, like in the case of *E. coli*, may be seen as a behavioural investment in the long-term maintenance of homeostasis given anticipated environmental changes. It is a form of anticipatory control allowing organisms to avoid deviations from homeostatic balance by engaging behaviours that enslave more flexible homeostatic mechanisms. As such, this example of anticipatory behaviour falls under the category of expected free energy-minimizing allostatic control.

Example 3 Allostatic Control in Slime Mould? Slime mould (*Physarum polycephalum*) is an amoeboid that can live as an individual cell or as a multicellular aggregate (a plasmodium). As plasmodia, the individual cells benefit from enhanced motility and more effective resource foraging (Dussutour et al. 2010). Plasmodia have continued to attract the interest and attention of biologists, cognitive scientists, and computer scientists over recent years because of their ability to display behaviour that is typically associated with organisms that have a brain (Vallverdú et al. 2018). For example, they have been shown to exhibit a form of habituation (Dussutour et al. 2010; Boisseau et al. 2016) and solve ‘the shortest path problem’ when placed in mazes (Nakagaki et al. 2000; Bonifaci et al. 2012). Importantly, *P. polycephalum* presents another fascinating example of anticipatory control in microorganisms. For example, Saigusa et al. (2008) have shown that plasmodia can anticipate and respond to periodic environmental changes. To do so, these researchers measured the locomotion speed of *Physarum* plasmodia when subjected to controlled periodic changes in two different experimental conditions: a “standard condition” at 26 °C and 90% humidity and a “dry stimulation” at 23 °C and 60% humidity. In the standard conditions, *Physarum* plasmodia migrated at a steady rate. When, however, subjected to periodic dry stimulation there is a drop in the plasmodium’s locomotion speed. Crucially, after three hourly instances of periodic exposure to the dry stimulation it was shown that plasmodia locomotion decreases spontaneously in standard conditions, coinciding with the time points (i.e., subsequent hours) where the next dry stimulations would have taken place. The existence of this kind of behaviour, Saigusa et al. suggest “means that the organism anticipates the next periodic environmental change” (2008, p. 4).

Why would such anticipatory behaviour be adaptively advantageous to *Physarum* plasmodia? Dry stimulation is physiologically challenging to *P. polycephalum*. In lasting dry conditions, plasmodia divert energy into formation of sclerotium (i.e., an encysted resting body that is resistant to desiccation) (Furuhashi 2002). One reasonable way of construing the adaptive nature of plasmodial locomotion-slowness is in terms of an effective trade-off of metabolic investment to secure the minimization of expected free energy. Less metabolic resources are expended in locomotion and more are diverted towards biochemical priming in preparation for the metabolically

expensive process of sclerotium formation. The trade-off roughly maps on to selecting between one of two policies: continued resource foraging and avoiding unexpected risky outcomes (i.e., drying out). Moreover, Saigusa et al. showed that after spontaneous locomotion-slowness ceases, the slowing down response *may be elicited again after only one exposure to dry stimulation*. This seems to strongly suggest to us that a plasmodium's anticipatory behaviour is the result of its having become a model of the periodic events in its environment over the course of its development.¹⁶ By investing its resources in prospective behaviour that corresponds to its learnt model of its 'act-on-able' environment, a plasmodium can avoid possible dyshomeostasis-inducing desiccation, thus maintaining its long-term homeostasis. The behaviour of *P. polycephalum* in other words exemplifies expected free energy minimizing allostatic control.

Example 4 Phenotypic Plasticity and Allostatic Control in Plants Unlike most animals and some microorganisms, plants are non-motile (i.e., sessile) organisms. Where motile organisms adapt to their changing environments by moving to bring about expected sensory states, adaptive behaviour in plants typically exploits phenotypic plasticity. Plasticity may be roughly defined as “the degree to which an organism can be changed in response to environmental signals” (Trewavas 2003, p. 12).¹⁷ It may take the form of either irreversible movement via *growth* (e.g., root nutation, circumnutation of helical organs, etc.) or reversible movement via increasing or decreasing *cellular turgor* (i.e., change in hydrostatic pressure) (Segundo-Ortín and Calvo 2019).¹⁸ Importantly, because such plastic behaviour requires time, “useful information, must be relevant to the future environment and conditions that the organism will eventually function in” (Novoplansky 2016, p. 59). As such, plasticity is an indicator of anticipation (Trewavas 2003). We will close this section by

¹⁶ Recall from Sect. 3 how Clark has expressed scepticism about Friston's talk of the organism becoming a model of its environment. We take the plasmodia example (and the other examples we discuss in this section) to show how it does indeed make sense to talk of the organism literally (and not only metaphorically) becoming a model of its environment in ontogenesis through epigenetic processes like phenotypic plasticity.

¹⁷ Phenotypic plasticity is a ubiquitous phenomenon in all living systems. For example, in neuronal organisms such plasticity is associated with the kind of remodelling of synaptic interactions that underlie memory formation and learning (Phillips 2006; Baluška and Levin 2016). The point we are making here is that phenotypic plasticity in plants, which plays a particularly crucial role in their ability, as sessile organisms, to actively adapt to the uncertainty of their external milieus. It should also be noted that the kind of development described by phenotypic plasticity is distinct from *unitary development* (i.e., the kind of development that occurs with maturation). While the latter may have adaptive advantages, it is not directed. Phenotypic plasticity can be directed and hence qualifies as an *adaptive behavioural strategy* (Trewavas 2003; Calvo and Keijzer 2011; Calvo and Friston 2017; Sims 2019).

¹⁸ The number of ways that a plant can plastically adapt to its changing environment may be easily conceptualized as corresponding to the number of ways animals and some microorganisms can alter their

looking at an example of anticipatory plant behaviour in purslane (*Portulaca oleracea*) seedlings.

In addition to water and minerals, light is an extremely important resource for plants. Photosynthetic processing by leaves allows plants to convert light energy into sugars which may be used to satisfy their metabolic requirements (Nevins 1995). Different species of plants however have different light requirements. Of those plants that thrive in more light, competition amongst them for this resource presents a significant kind of selection pressure. These kinds of “sun-loving plants” that are found in open habitats, when shaded by other vegetation, respond by elongating their stems, increasing their chances of escaping the shade of their neighbouring sun-loving competitors and simultaneously casting shade upon them (Aarssen 1995; Novoplansky 2016). More generally, plants’ sensitivity to the ratio of red light (wavelength 650–700 nm) and to far-red light (wavelength 700–750 nm) allows them to respond to competing plants, elongating both their stems and leaves in addition to reorienting their leaves to a more optimal position for photosynthetic light exposure (Smith 2000). When the proportion of red light is lower than that of far-red light this ratio may be used as a cue for the presence of competing plants. This is because inanimate objects (e.g., dead plants) reduce a full range of spectral light frequencies, whereas plants absorb red and reflect far-red light, leaving a low red/far-red light ratio as a proxy for the presence of shading competitors (Casal and Smith 1989; Novoplansky 1990, 1991, 2016).

The fact that low red/far-red light can be used by plants to predict the presence of competing neighbours was the basis of a revealing experiment by Novoplansky et al. (1990). These researchers were able to demonstrate that the seedlings of the common purslane (*Portulaca oleracea*) not only detect spectral changes as cues for the presence and the direction of neighbouring plants but they also engage in anticipatory behaviour. They reorient their growth *prior to being shaded by their neighbours*. Novoplansky and colleagues showed that young purslane seedlings engage in such anticipatory behaviour by using a process they call “horizon filtering”. They used shallow (2 cm) plastic rim filters to simulate the kinds of spectral effects that small neighbouring plants would have on seedlings. These rims featured two separate coloured filters, a grey and green filter respectively. Whereas the grey filter simulated neutral shade (i.e., akin to the shade of an inanimate object), the green filter reduced red/far-red light frequencies (akin to the shade of neighbouring seedling plants). In each experimental treatment, the plastic rim was placed around the seedling so that the green filter was located in a different cardinal direction, leaving the grey filter surrounding the rest of the seedling.

Footnote 18 (continued)

spatial relationship to their environments through quick and motility-driven behaviours. This point is expressed elegantly by Trewavas when he writes: “Just as animal behaviour is constrained by genetic capabilities, so ultimate genetic constraints on phenotypic change will be present. But with plants refining their discrimination to local conditions, perhaps the enormous numbers of distinguishable phenotypes correspond well with the number of behavioural variations available to any animal” (2003, p.13). For a more recent comparison of plant and animal behaviour see Calvo and Trewavas (2020).

Novoplansky and colleagues found that the seedlings grew towards grey but away from the green despite the fact that the filter's shallow height only filtered light with respect to the position of the seeding at sunrise and sunset. In other words, despite being exposed to photosynthetic light for most of the day the seedlings invested their metabolic resources in prospectively growing away from the direction in which *future shading by neighbouring plants was expected*. These results may be taken to suggest that “rather than opportunistically maximise present absorption of photosynthetic light, [purslane seedlings] are able to perceive and integrate directional spectral cues and respond in ways that maximize their expected total long-term light absorption and growth” (Novoplansky et al. 2016, p. 62–63). Given that future photosynthetic light depletion for a sun-loving plant means eventual dyshomeostasis, we take the anticipatory behaviour exhibited by *P. oleracea* seedlings as an example of behaviour that aims at outcomes expected to minimise future free energy.¹⁹

We take these four examples to provide compelling evidence for two claims. First these organisms become a model of the opportunities and challenges of their environmental conditions. The developmental processes by which organisms develop an organisation that models their environment can take various forms and can occur across different timescales. Importantly, instantiating an internal dynamics that models correlations and associations between cues is only possible when the environment contains regular structure that may be capitalized upon. Second, the model these organisms develop is used for the purpose of allostatic control. Organisms expend energy and expose themselves to stressful conditions that temporarily incur an increase in free energy but they do so with the aim of minimising expected free energy. Should we conclude then that organisms like the ones we have discussed that lack a nervous system can nevertheless cognise? We have argued that the FEP supports allostatic control as the mark of cognition. By this criteria, we should indeed conclude that evolutionarily primitive organisms are capable of cognition.

Conclusion

We have shown how the FEP provides a mark of the cognitive that if it is accepted would warrant the ascription of cognitive processes to organisms that lack a nervous system. A distinction is commonly made within cognitive science between different forms of sensorimotor control—Pavlovian, habit-based and goal-directed. Such a distinction would seem to support a picture of evolution in which the earliest forms of life relied upon a sensorimotor architecture of homeostatic control, and gradually forms of instrumental began to evolve. Goal-directed control as it occurs in humans evolved out of these earlier forms of instrumental control. Organisms started out with a narrow and fixed repertoire of behaviours but as they developed instrumental

¹⁹ For a more recent example where young pea plants have been shown to detect and integrate information about their environmental resources to guide anticipatory behaviour with respect to future growth conditions see Shemesh et al. (2010).

forms of control, so their behavioural repertoire expanded. The FEP is sometimes taken to support this picture of the evolution of cognition (see e.g. the contribution of Corcoran and colleagues in this volume and Clark 2016). We've argued however that FEP actually complicates this picture in productive ways suggesting that even in Pavlovian learning and through processes of phenotypic plasticity the organism is proactively anticipating the implications of its actions for future free energy and selecting those actions that minimise expected free energy. We take the FEP to support a picture of complex adaptive systems as goal-directed and purposeful in their behaviours. Even organisms that lack a nervous system exhibit primitive agency acting with the long-term purpose of keeping future free energy to a minimum. Their behaviour is not reactive but is anticipatory. These organisms seem to be capable of selecting actions based on their anticipated implications for free energy. Thus they challenge a picture of sensorimotor control in which instrumental forms of control that take into account the outcomes of action evolved out of a purely Pavlovian and reflexive control. Instrumental forms of allostatic control may have already been present at the beginning of the evolution of life. We have argued that allostatic control—an affect-based form of regulation—may prove to be the mark of the cognitive. Thus it is an implication of our paper that cognition may have already been present at the very beginning of the evolution of life.

Acknowledgements JK's work on this paper was supported by an Amsterdam Brain and Cognition research grant and by a European Research Council grant (no 679190) awarded to Erik Rietveld. We are grateful for helpful comments from two anonymous reviewers and from Dave Ward.

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

- Aarssen LW (1995) Hypothesis for the evolution of apical dominance in plants: implications for the interpretation of overcompensation. *Oikos* 74:149–156
- Adams F (2018) Cognition wars. *Stud Hist Philos Sci* 68:20–30
- Adams, F. (2019). The elusive extended mind: extended information-processing doesn't equal extended mind. In M. Colombo, E. Irvine, & M. Stapleton (Ed's) *Andy Clark and his Critics*. Oxford, UK: Oxford University Press, pp.21–31.
- Alfieri N (2008) Cellular stress response: from homeostatic to allostatic perspective. *Biosci Hypotheses* 1:330–331
- Allen C (2017) On (not) defining cognition. *Synthese* 194(11):4233–4249
- Anderson M (2017) Of Bayes and bullets: an embodied, situated, targeting-based account of predictive processing. *Philosophy and predictive processing*. MIND Group, Frankfurt am Main
- Andrews, M. (2020). The math is not the territory: navigating the Free Energy Principle. Preprint, downloaded from: <http://philsci-archive.pitt.edu/18315/>. Accessed 10 Jan 2021
- Auletta G (2013) Information and metabolism in bacterial chemotaxis. *Entropy* 15(1):311–326

- Baltieri M, Buckley CL, Bruineberg J (2020) Predictions in the eye of the beholder: an active inference account of Watt governors
- Baluška F, Levin M (2016) On having no head: cognition throughout biological systems. *Front Psychol* 7:902. <https://doi.org/10.3389/fpsyg.2016.00902>
- Baluška F, Mancuso S (2013) Root apex transition zone as oscillatory zone. *Front Plant Sci* 4:354. <https://doi.org/10.3389/fpls.2013.00354>
- Barrett L (2011) *Beyond the brain: how the body and the environment shape animal and human minds*. Princeton University Press, New Jersey
- Boisseau RP, Vogel D, Dussutour A (2016) Habituation in non-neural organisms: evidence from slime moulds. *Proc Biol Sci* 283(1829):20160446. <https://doi.org/10.1098/rspb.2016.0446>
- Bonifaci V, Mehlhorn K, Varma G (2012) Physarum can compute shortest paths. *J Theor Biol* 309:121–133
- Bruineberg J, Kiverstein J, Rietveld E (2018) The anticipating brain is not a scientist: the free-energy principle from an ecological-enactive perspective. *Synthese* 195:2417–2444
- Bruineberg J, Rietveld E, Parr T, Van Maanen L, Friston KJ (2018) Free-energy minimisation in joint agent-environment system: a niche-construction perspective. *J Theor Biol* 455:161–178
- Bruineberg J, Dolega K, Dewhurst J, Baltieri M (2020) The emperor's new Markov blankets. Preprint, downloaded from <http://philsci-archive.pitt.edu/18467/> Accessed 15 Jan 2021
- Burge T (2009) Primitive agency and natural norms. *Philos Phenomenol Res* 79(2):251–278
- Burge T (2010) *The origins of objectivity*. Oxford University Press, Oxford
- Calvo P, Friston K (2017) Predicting green: really radical (plant) predictive processing. *J R Soc Interface* 14:20170096. <https://doi.org/10.1098/rsif.2017.0096>
- Calvo P, Trewavas A (2020) Cognition and intelligence of green plants. *Information for animal scientists*. *Biochem Biophys Res Commun*. <https://doi.org/10.1016/j.bbrc.2020.07.139>
- Calvo Garzon P, Keijzer F (2011) Plants: adaptive behavior, root brains and minimal cognition. *Adapt Behav* 19:155–171
- Campbell JO (2016) Universal Darwinism as a process of Bayesian inference. *Front Syst Neurosci* 10:49
- Cannon WB (1929) Organization for physiological homeostasis. *Physiol Rev* 9(3):399–431
- Casal JJ, Smith H (1989) The function, action and adaptive significance of photochrome in light-grown plants. *Plant Cell Environ* 12:855–862
- Cisek P (1999) Beyond the computer metaphor: behaviour as interaction. *J Conscious Stud* 6(11–12):125–142
- Clark A (2011) Finding the mind. *Philos Stud* 152(3):447–461
- Clark A (2016) *Surfing uncertainty: prediction, action, and the embodied mind*. Oxford University Press, New York
- Clark A (2017) How to knit your own markov blanket. In: Metzinger TK, Wiese W (eds) *Philosophy and predictive processing*. MIND Group, Frankfurt am Main, pp 34–59
- Clark A (2019) Replies to critics: in search of the embodied, extended, enactive, predictive (EEE-P) mind. In: Colombo M, Irvine E, Stapleton M (eds) *Andy Clark and his Critics*. Oxford University Press, Oxford, pp 266–303
- Colombo M, Seriès P (2012) Bayes in the brain: on Bayesian modelling in neuroscience. *Br J Philos Sci* 63:697–723
- Colombo M, Wright C (2018) First principles in the life sciences: the free-energy principle, organicism, and mechanism. *First Principles in Science, Synthese*
- Colombo M, Elkin L, Hartmann S (2018) Being realist about Bayes and the predictive processing theory of mind. *Br J Philos Sci*. <https://doi.org/10.1093/bjps/axy059>
- Conant R, Ashby WR (1970) Every good regulator of a system must be a model of that system. *Int J Syst Sci* 1(2):89–97
- Corcoran A, Pezzulo G, Hohwy J (2020) From allostatic agents to counterfactual cognisers: active inference, biological regulation, and the origins of cognition. *Biol Philos*. <https://doi.org/10.1007/s10539-020-09746-2>
- Craik KJW (1943) *The nature of explanation*. Cambridge University Press, Cambridge
- Davis KAJ (2016) Adaptive homeostasis. *Mol Aspects Med*. <https://doi.org/10.1016/j.mam.2016.04.007>
- Dennett D (1995) *Darwin's dangerous idea. evolution and the meaning of life*. Simon & Schuster, New York
- Di Paolo EA, Buhrmann T, Barandiaran X (2017) *Sensorimotor life: an enactive proposal*. Oxford University Press, Oxford

- Dussutour A, Latty T, Beekman M, Simpson SJ (2010) Amoeboid organism solves complex nutritional challenges. *Proc Natl Acad Sci USA* 107:4607–4611. <https://doi.org/10.1073/pnas.0912198107>
- Friston KJ (2010) The free-energy principle: a unified brain theory? *Nat Rev Neurosci* 11(2):127–138
- Friston KJ (2011) Embodied inference: Or “I think therefore I am, if I am what I think.” In: Tschacher W, Bergomi C (eds) *The Implications of Embodiment: Cognition and Communication*. Exeter, Imprint Academic, pp 89–125
- Friston KJ (2012) A free energy principle for biological systems. *Entropy* 14(11):2100–2121
- Friston KJ (2013) Life as we know it. *J R Soc Interface*. <https://doi.org/10.1098/rsif.2013.0475>
- Friston KJ, Stephan KE (2007) Free-energy and the brain. *Synthese* 159(3):417–458
- Friston KJ, Rigoli F, Ognibene D, Mathys C, Fitzgerald T, Pezzulo G (2015) Active inference and epistemic value. *Cogn Neurosci* 6(4):187–214
- Fulda F (2017) Natural agency: the case of bacterial cognition. *J Am Philos Assoc* 3(1):69–90
- Fultot M, Frazier PA, Turvey MT, Carello C (2019) What are nervous systems for? *Ecol Psychol* 31(3):218–234
- Furuhashi K (2002) Identification and characterization of a cathepsin B-like protease in *Physarum sclerotium*. *Int J Biochem Cell Biol* 34(10):1308–1316. [https://doi.org/10.1016/s1357-2725\(02\)00072-9](https://doi.org/10.1016/s1357-2725(02)00072-9) (PMID: 12127582)
- Gagliano M, Vyazovskiy VV, Borbély AA, Grimonprez M, Depczynski M (2016) Learning by association in plants. *Sci Rep* 6:38427. <https://doi.org/10.1038/srep38427>
- Godfrey-Smith P (2002) Environmental complexity and the evolution of cognition. In: Sternberg RJ, Kaufman JC (eds) *The Evolution of Intelligence*. Psychology Press, London, pp 233–249
- Godfrey-Smith P (2016) Individuality, subjectivity, and minimal cognition. *Biol Philos* 31(6):775–796
- Goodson FE (2003) *The evolution and function of cognition*. Lawrence Erlbaum Associates Inc., Mahwah
- Hellingwerf KJ (2005) Bacterial observations: a rudimentary form of intelligence? *Trends Microbiol* 13:152–158. <https://doi.org/10.1016/j.tim.2005.02.001>
- Hohwy J (2013) *The Predictive Mind*. Oxford University Press, Oxford
- Hohwy J (2015) The neural organ explains the mind. In: Metzinger T, Windt JM (eds) *Open MIND*. MIND Group, Frankfurt am Main, pp 1–23
- Hohwy J (2019) Quick ‘n’ lean or slow and rich? Andy Clark on predictive processing and embodied cognition. In: Colombo M, Irvine E, Stapleton M (eds) *Andy Clark and his Critics*. Oxford University Press, Oxford, pp 191–205
- Hohwy J (2020) Self-supervision, normativity and the free energy principle. *Synthese*. <https://doi.org/10.1007/s11229-020-02622-2>
- Hurley S (2010) The varieties of externalism. In: Menary R (ed) *The Extended Mind*. MIT Press, Cambridge, pp 101–154
- Hutto D, Myin E (2017) *Evolving enactivism: basic minds without content*. MIT Press, Cambridge
- Keijzer FA (2020) Demarcating cognition: the cognitive life sciences. *Synthese* 198:137–157
- Kiefer A, Hohwy J (2017) Content and misrepresentation in hierarchical generative models. *Synthese* 1–29
- Kirchhoff M, Froese T (2017) Where there is life there is mind: In: Support of a Strong Life-Mind Continuity Thesis. *Entropy*, 19(4), p. 169
- Kirchhoff M, Parr T, Palacios E, Friston K, Kiverstein J (2018) The Markov blankets of life: autonomy, active inference and the free energy principle. *J R Soc Interface* 15(138):20170792
- Kiverstein J (2020) Embodied cognition and the neural reuse hypothesis. In: Lerner AJ, Cullen S, Leslie SJ (eds) *Current controversies in philosophy of cognitive science*. Routledge, London
- Linson A, Calvo P (2020) Zoocentrism in the weeds? Cultivating plant models for cognitive yield. *Biol Philos* 35:49
- Lyon P (2006) The biogenic approach to cognition. *Cogn Process* 7(1):11–29
- Lyon P (2015) The cognitive cell: bacterial behavior reconsidered. *Front Microbiol* 6:264
- Maris AF, Assumpção AL, Bonatto D, Brendel M, Henriques JA (2001) Dauxic shift induced stress resistance against hydroperoxides in *Saccharomyces cerevisiae* is not an adaptive stress response and does not depend upon mitochondria. *Curr Genet* 39:137–149
- Mitchell A, Romano GH, Groisman B, Yona A, Dekel E, Kupiec M, Dahan O, Pilpel Y (2009) Adaptive prediction of environmental changes by microorganisms. *Nature* 460(7252):220–224
- Morgan CL (1894) *An introduction to comparative psychology*. Walter Scott, London

- Müller BS, Di Primio F, Lengeler JW (2001) Contributions of minimal cognition to flexibility. In: SCI 2001 Proceedings of the 5th World Multi-Conference on Systemics, Cybernetics and Informatics, vol. 15, pp. 93–98
- Nakagaki T, Yamada H, Tóth A (2000) Maze-solving by an amoeboid organism. *Nature* 407:470
- Nevins DJ (1995) Sugars: their origin in photosynthesis and subsequent biological interconversions. *Am J Clin Nutr* 61(4):915S–921S. <https://doi.org/10.1093/ajcn/61.4.915S> (PMID: 7900689)
- Noë A (2004) *Action in Perception*. MIT Press, Cambridge
- Novoplansky A (1991) Developmental responses of portulaca seedlings to conflicting spectral signals. *Oecologia* 88:138–140
- Novoplansky A (2016) Future perception in plants. In: Nadin M (ed) *Anticipation across disciplines*. Springer, Berlin, pp 57–70
- Novoplansky A, Cohen D, Sachs T (1990) How portulaca seedlings avoid their neighbors. *Oecologia* 82:490–493
- Pfeifer R, Bongard J (2007) *How the body shapes the way we think: a new view of intelligence*. MIT Press, Cambridge
- Phillips K (2006) Phenotypic plasticity. *J Exp Biol* 209:1–3
- Pinto D, Mascher T (2016) (Actino) Bacterial “intelligence”: using comparative genomics to unravel the information processing capacities of microbes. *Curr Genet* 62:487–498. <https://doi.org/10.1007/s00294-016-0569-3>
- Ramstead MJ, Badcock PB, Friston KJ (2018) Answering Schrödinger’s question: a free-energy formulation. *Phys Life Rev* 24:1–16
- Ramstead MJ, Kirchhoff MD, Friston KJ (2019) A tale of two densities: active inference is enactive inference. *Adapt Behav*. <https://doi.org/10.1177/1059712319862774>
- Ramstead M, Friston KJ, Hippolito I (2020b) Is the free-energy principle a formal theory of semantics? From variational density dynamics to neural and phenotypic representations. *Entropy* 22(8):889
- Ramstead M, Hesp C, Tschantz A, Smith R, Constant A, Friston KJ (2020a) Neural and phenotypic representation under the free-energy principle
- Reid CR, Latty T (2016) Collective behaviour and swarm intelligence in slime moulds. *FEMS Microbiol Rev* 40(6):798–806
- Rosen R (1985/2012) *Anticipatory systems: philosophical, mathematical, and methodological foundations*. Pergamon, Oxford.
- Saigusa T, Tero A, Nakagaki T, Kuramoto Y (2008) Amoebae anticipate periodic events. *Phys Rev Lett* 100:18101
- Salek M, Carrara F, Vicente F, Gausto J, Stocker R (2019) Bacterial chemotaxis in a microfluidic T-maze reveals strong phenotypic heterogeneity in chemotactic sensitivity. *Nat Commun* 10:1877. <https://doi.org/10.1038/s41467-019-09521-2>
- Schwartenbeck P, Passecker J, Hauser TU, FitzGerald TH, Kronbichler M, Friston KJ (2019) Computational mechanisms of curiosity and goal-directed exploration. *eLife* 8:e41703
- Segundo-Ortín M, Calvo P (2019) Are plants cognitive? A reply to Adams. *Studies in history and philosophy of science. Stud Hist Philos Sci Part A*. <https://doi.org/10.1016/j.shpsa.2018.12.001>
- Seth AK (2015) The cybernetic bayesian brain: from interoceptive inference to sensorimotor contingencies. In: Metzinger T, Windt JM (eds) *Open MIND*. MIND Group, Frankfurt am Main, pp 1–24
- Seth AK (2019) Being a beast machine: the origins of selfhood in control-oriented interoceptive inference. In: Colombo M, Irvine E, Stapleton M (eds) *Andy Clark and his Critics*. Oxford University Press, Oxford, pp 238–253
- Shapiro JA (2007) Bacteria are small but not stupid: cognition, natural genetic engineering and sociobacteriology. *Stud Hist Philos Biol Biomed Sci* 38(4):807–819
- Shemesh H, Arbib A, Gersani M, Ovadia O, Novoplansky A (2010) The effects of nutrient dynamics on root patch choice. *PLoS ONE* 5:e10824
- Sims M (2019) Minimal-perception: responding to the challenges of perceptual constancy and veridicality with plants. *Philos Psychol*. <https://doi.org/10.1080/09515089.2019.1646898>
- Sims M (2020) How to count biological minds: symbiosis, the free-energy principle, and reciprocal multiscale integration (2020). *Synthese*. <https://doi.org/10.1007/s11229-020-02876-w>
- Smith H (2000) Phytochromes and light signal perception by plants—an emerging synthesis. *Nature* 407:585–591
- Stephan KE, Manjaly ZM, Mathys CD, Weber LAE, Paliwal S, Gard T, Tittgemeyer M, Fleming SM, Haker H, Seth AK, Petzschner FH (2016) Allostatic self-efficacy: a metacognitive theory of dyshomeostasis-induced fatigue and depression. *Front Hum Neurosci* 10(550):1–27

- Sterling P (2012) Allostasis: a model of predictive regulation. *Physiol Behav* 106(1):5–15
- Stich S (1990) *The Fragmentation of Reason: Preface to a Pragmatic Theory of Cognitive Evaluation*. MIT Press, Cambridge, MA
- Tagkopoulos I, Liu Y-C, Tavazoie S (2008) Predictive behavior within microbial genetic networks. *Science* 320(5881):1313–1317
- Trewavas A (2003) Aspects of plant intelligence. *Ann Bot* 92:1–20
- Trewavas A (2014) *Plant behaviour and intelligence*. Oxford University Press, Oxford
- Tribus M (1961) *Thermodynamics and thermostatics: an introduction to energy, information and states of matter with engineering applications*. D. Van Nostrand, New York
- Tschantz A, Seth AK, Buckley CL (2020) Learning action-oriented models through active inference. *PLoS Comput Biol* 16(4):e1007805
- Vallverdú J, Castro O, Mayne R, Talanov M, Levin M, Baluška F, Gunji Y, Dussutour A, Zenil H, Adamtzky A (2018) Sime mould: the fundamental mechanisms of biological cognition. *Biosystems* 165:57–70
- Van Es T (2020) Living models or life modelled? On the use of models in the free energy principle. *Adapt Behav*. <https://doi.org/10.1177/1059712320918678>
- Van Duijn M, Keijzer FA, Franken D (2006) Principles of minimal cognition: Casting cognition as sensorimotor coordination. *Adaptive Behavior* 14(2):157–170
- Varela F, Thompson E, Rosch E (1991) *The embodied mind: cognitive science and human experience*. MIT Press, Cambridge
- Walsh DM (2015) *Organisms, agency and evolution*. Cambridge University Press, Cambridge
- Wang J, Atolia E, Hua B, Savir Y, Escalante-Chong R, Springer M (2015) Natural variation in preparation for nutrient depletion reveals a cost-benefit tradeoff. *PLoS Biol* 13(1):e1002041. <https://doi.org/10.1371/journal.pbio.1002041>
- Wiese W, Metzinger T (2017) Vanilla predictive processing for philosophers: a primer on predictive processing. In: Metzinger T, Wiese W (eds) *Philosophy and predictive processing*. MIND Group, Frankfurt am Main
- Wolf DM, Fontaine-Bodin L, Bischofs I, Price G, Keasling J et al (2008) Memory in microbes: quantifying history-dependent behavior in a bacterium. *PLoS ONE* 3(2):e1700. <https://doi.org/10.1371/journal.pone.0001700>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.