

The feeling of grip: novelty, error dynamics, and the predictive brain

Julian Kiverstein¹  · Mark Miller² ·
Erik Rietveld¹ 

Received: 27 April 2017 / Accepted: 30 September 2017 / Published online: 23 October 2017
© The Author(s) 2017. This article is an open access publication

Abstract According to the free energy principle biological agents resist a tendency to disorder in their interactions with a dynamically changing environment by keeping themselves in sensory and physiological states that are expected given their embodiment and the niche they inhabit (Friston in *Nat Rev Neurosci* 11(2):127–138, 2010. doi:[10.1038/nrn2787](https://doi.org/10.1038/nrn2787)). Why would a biological agent that aims at minimising uncertainty in its encounters with the world ever be motivated to seek out novelty? Novelty for such an agent would arrive in the form of sensory and physiological states that are unexpected. Such an agent ought therefore to avoid novel and surprising interactions with the world one might think. Yet humans and many other animals find play and other forms of novelty-seeking and exploration hugely rewarding. How can this be understood in frameworks for studying the mind that emphasise prediction error min-

Many thanks to Andy Clark for excellent feedback and suggestions on earlier drafts of this paper, and to Jelle Bruineberg for helping us to understand the recent literature in cognitive neuroscience on epistemic value. We are extremely grateful for detailed reviews we received from two anonymous reviewers that helped us to significantly improve and sharpen our arguments. Erik Rietveld would like to acknowledge the research funding awarded by the Netherlands Organisation for Scientific Research (NWO) in the form of a VIDI-grant and the European Research Council in the form of ERC Starting Grant 679190 (EU Horizon 2020) for the project AFFORDS-HIGHER.

✉ Julian Kiverstein
j.d.kiverstein@amc.uva.nl
Erik Rietveld
d.w.rietveld@amc.uva.nl
<http://www.erikrietveld.com>

¹ Academic Medical Centre, Department of Philosophy, ILLC, University of Amsterdam, Amsterdam, The Netherlands

² School of Philosophy, Psychology and Language Sciences, University of Edinburgh, Edinburgh, UK

imisation? This problem has been taken up in recent research concerned with epistemic action—actions an agent engages in to reduce uncertainty. However that work leaves two questions unanswered, which it is the aim of our paper to address. First, no account has been given yet of why it should feel good to the agent to engage the world playfully and with curiosity. Second an appeal is made to precision-estimation to explain epistemic action, yet it remains unclear how precision-weighting works in action more generally, or active inference. We argue that an answer to both questions may lie in the bodily states of an agent that track the rate at which free energy is being reduced. The recent literature on the predictive brain has connected the valence of emotional experiences to the rate of change in the reduction of prediction error (Joffily and Coricelli in *PLoS Comput Biol* 9(6):e1003094, 2013. doi:10.1371/journal.pcbi.1003094; Van de Cruys, in Metzinger and Wiese (eds) *Philosophy and predictive processing*, vol 24, MIND Group, Frankfurt am Main, 2017. doi:10.15502/9783958573253). In this literature valenced emotional experiences are hypothesised to be identical with changes in the rate at which prediction error is reduced. Experiences are negatively valenced when overall prediction error increases and are positively valenced when the sum of prediction errors decrease. We offer an ecological-enactive interpretation of the concept of valence and its connection to rate of change of prediction error. We show how rate of change should be understood in terms of embodied states of affordance-related action readiness. We then go on to apply this ecological-enactive account of error dynamics to provide an answer to the first question we have raised: It may explain why it should feel good to an agent to be curious and playful. Our ecological-enactive account also allows us to show how error dynamics may provide an answer to the second question we have raised regarding how precision-weighting works in active inference. An agent that is sensitive to rates of error reduction can tune precision on the fly. We show how this ability to tune precision on the go can allow agents to develop skills for adapting better and better to the unexpected, and search out opportunities for resolving uncertainty and progressing in its learning.

Keywords Free energy principle · Predictive processing · Epistemic value · Exploit-explore dilemma · Curiosity · Error dynamics · Emotion · Decision-making · Embodied cognition · Enactive cognitive science

Introduction

Cognitive neuroscience is on the brink of formulating an elegant unifying theory that shows how the principles that define living systems, also explain the workings of the human mind. The foundations of this theory come from a mathematically complex principle—the so-called “free energy principle” (FEP), which can be applied to every biological system that resists a tendency to disorder (Friston 2009, 2010, 2013; Kirchoff and Froese 2017).¹ Friston has proposed that everything that can change in the

¹ Elsewhere we have provided a detailed philosophical overview and analysis of the free-energy principle (FEP), showing how the FEP should be interpreted in ecological-enactive terms (Bruineberg and Rietveld 2014; Bruineberg et al. 2016). Here we restrict our focus to dealing with a problem that seems to arise for FEP when it comes to accounting for behaviours that are motivated by curiosity, exploration and playfulness.

brain will change so as to maintain the adaptive fit of the agent to its dynamically changing environment. The brain (as an integrated part of larger agent-environment system) should steer the agent's interactions with the world so as to maximize the probability that it stays in the physiological and sensory states that are expected given its embodiment and the niche it inhabits. For example, the human body has a high probability of having a temperature of around 37 °C. Homeostatic processes in the brain should then regulate body temperature so that the thermal states of the body stay as close as possible to this expected value. In other words, the brain should be organized in such a way as to suppress “surprise”, which will remain low when the organism maintains itself in physiological and sensory states that are expected, and will increase should the organism find itself in states that are improbable and hence unexpected. “Surprise” is a technical term and relates to predictions of sensory and physiological states over time—it is future oriented. More precisely, surprise is associated with trajectories or sequences of sensory input; thereby lending it a dynamic and anticipatory aspect. In this treatment, we will be concerned with the surprise of extended sensory outcomes, consequent upon the pursuit of action policies. The brain contributes to ensuring that the agent avoids surprise by anticipating how the agent's sensory and physiological states will change over time as it moves through its environment. So long as the brain succeeds in minimizing the divergence between the change in sensory states that it anticipates and the changes in sensory states that actually ensue, it will succeed in keeping the agent away from surprising outcomes, and maintain the agent's adaptive fit to its environment.

Friston provides a precise mathematical framework for quantifying the value of this divergence between the change in sensory states the brain anticipates and the change that actually occurs, using the information-theoretic concept of free-energy. Friston claims that self-organising adaptive systems will avoid surprising states by having a functional organization that continuously minimizes free-energy over the long run.² Free energy is related to entropy, it is a measure of the biological agent's order. In thermodynamics and statistical mechanics it refers to the amount of energy that can be extracted from a system and put to work (McEvoy 2002; Clark 2013: p. 186), which is roughly the “difference between the energy and the entropy of the system” (Friston and Stephan 2007: p. 419). The concept of free energy at work in the free energy principle is variational free energy, a “measure of statistical probability distributions” (Friston and Stephan 2007: p. 420). More precisely, what free energy measures is the divergence between a probability distribution typically interpreted as encoding “prior beliefs” about the hidden statistical structure in data, and current sensory evidence. This divergence provides a means of quantifying the information that is available for use in the current sensory evidence. The lower the free energy, the better the system's “beliefs”. This is to say that the system's cognitive resources are being put to work

Footnote 1 continued

We provide only as much theoretical background in this paper as is needed for generating the problem. For readers interested in learning more about FEP we recommend they consult the papers by Friston we cite here and the recent special issue on predictive brains in this journal.

² When these dynamics are instantiated in an embodied organism the natural result of this ongoing reduction of error at all levels is the generation and maintenance of homeostasis (Seth 2015; Bruineberg et al. 2016).

in ways that are maximally useful for adapting it to the environment. Free energy increases when the biological agent finds itself in states (potentially life-threatening) that are unexpected relative to its beliefs about the world. The more free energy, which is to say the more often the biological agent finds itself in unexpected sensory and physiological states, the less useful work the biological agent's "beliefs" about the world do.³

Wherever there is free energy there is room for improvement—there is sensory prediction error that is not currently accommodated by one's model. This prediction error can then be accommodated through action, or by improving one's existing model. In this paper we will be concerned with the prediction of the temporally-extended sensory consequences of action, or expected free energy. We say more about the latter concept below. The FEP therefore claims that biological systems are organised in such a way as to minimise free energy continuously over time, which is equivalent to minimising uncertainty in an agent's active engagement with its environment.

If free-energy minimisation is a fundamental organising principle of the brain why is it then that our brains don't steer us towards environments in which sensory states can be easily predicted such as empty dark rooms? In the next section we outline the dark room problem and the solution that has been proposed. We agree with others that the dark room problem is in some ways a red herring (e.g. [Friston et al. 2012b](#)), however we will argue that a significant problem remains. The real problem is that of explaining why a biological system that acted based on the imperative to resist a tendency to disorder would be curious, motivated to explore its environment and seek out novelty. The free energy principle would seem to imply that valuable states are the ones the agent expects to be in. Yet curiosity and playfulness will more often than not lead an agent into states that are unexpected. Thus it looks at first glance as if a free energy minimising (FEM) agent ought not to be a curious and playful agent.

In Sect. 1 we outline this challenge in a little more detail. Section 2 shows how the problem has been addressed in recent work on "epistemic action". Epistemic actions are actions an agent engages in to reduce uncertainty. They allow the agent to "disclose information" through exploration "that enables pragmatic actions in the long run" ([Friston et al. 2015](#): p. 2).⁴ The research on epistemic action leaves two questions unanswered. First, it fails to explain why it feels good to the agent to engage the world playfully and with curiosity. Pleasure is a part of the value of curiosity and play for agents like us. Existing accounts of epistemic action do not fully explain how value works in epistemic action. Second, an appeal is made to precision-estimation

³ Free energy is for this reason sometimes defined in terms of accuracy minus complexity ([Friston 2010](#)). Accuracy is a function of how much prediction error the model produces over time—how well does the model do at reducing surprisal over the long-term. Complexity is measured technically in terms of Kubler–Leibler (K–L) divergence. It refers to the divergence between the prior probability of the hypothesis and the hypothesis selected based on the evidence. Complexity is high when many changes were made to the priors to fit the evidence (i.e. the divergence is large). Complexity leads to models that may fit current evidence very well but end up doing worse over time, generalising poorly to new situations (they are overfitted) ([Hohwy 2015](#): p. 5).

⁴ Pragmatic actions are actions that minimise free energy directly leading the agent to occupy the states they expect to be in based on past learning. We discuss Friston and colleagues' distinction between pragmatic and epistemic actions in more detail in the next section.

to explain epistemic action, yet it remains unclear how precision-weighting works in active inference. We argue that the answer to both questions may be found in the bodily states of an agent that track the rate at which free energy is being reduced. In Sect. 3 we take up the first of these questions and show how the agent can be sensitive to rates of free energy minimisation (FEM). This information about rate of change is given corporeally as states of affordance-related action readiness that are simultaneously affective and behavioural (Bruineberg & Rietveld, 2014; Rietveld, Denys & van Westen, 2017). We show how felt states of action readiness can account for the positive and negative hedonic tone that is often a feature of novel experience. In Sect. 4 we turn to the second question about precision-weighting and show how sensitivity to rate of change may play a role in tuning precision on the fly. This can ensure that the agent is steered towards opportunities for reducing uncertainty. We finish up in Sect. 5 by showing how an agent that is sensitive to error dynamics (rate of FEM) will be a curious agent, motivated to explore and play in its environment.

1 Worries about dark rooms

In FEP, agents act so as to keep themselves within expected sensory states given their embodiment and the niche they inhabit. The value of a sensory state is a function of how surprising it is. We are using the term “surprise” here in the technical sense introduced above that makes a conceptual connection between “surprise” and sensory states that are highly probable and thus expected given an agent’s embodiment and the niche it lives in. Unsurprising states (states highly frequented) are expected and are thus highly valued. Surprising states are not expected (they are improbable), and thus negatively valued. Positively valued states are often associated with reward (sometimes understood in terms of pleasure), while negatively valued states are typically associated with punishment and are consequently aversive. It follows that unsurprising states should be associated with pleasure, according to FEP and surprising states should be aversive (Friston et al. 2012a). A novel outcome of this perspective is that while it might feel as though we seek pleasures and avoid pains and so end up frequenting pleasurable states more often than painful states, according to FEP *highly frequented states are themselves the rewards*. Through a process known as “active inference” the agent acts to keep itself in states that are expected. A consequence of minimising free energy is that some states are occupied more than others. These are the states that are positively valued by an agent (Friston et al. 2014: p. 2). This means that we are not so much drawn to rewards, but are instead rewarded for reducing errors between expected and actual states. Traditional reinforcement learning models describe goal-directed behaviour as a product of the agent working out how best to maximize an expected reward (Schutz et al. 1997; Sutton and Barto 1998). In active inference the rewarding states are the states the agent learns to expect to occupy through a process of approximate Bayesian inference. Subsequent behaviour unfolds as the system attempts to reduce the discrepancies between the current state and the expected reward state (Schwartenbeck et al. 2014).

This view of decision-making as the outcome of active inference comes with a puzzle. If highly frequented (more expected) states are themselves rewarding and less frequented states (more uncertainty) are aversive, why then should agents ever be

motivated to seek out novelty? A good strategy for guaranteeing that one remains in sensory states that are expected might seem to be to seek out a simple, static environment such as a dark empty room in which very little ever changes (Friston et al. 2012b). The agent adopting such a strategy would be pretty much guaranteed to only occupy the sensory states they expected to occupy given their model of the dark room. Nothing unexpected happens in a dark empty room. Thus, once one has learned a good model of such an environment, one is pretty much guaranteed to remain in the states one expects.

The problem of why free energy minimising agents tend not to retreat and hide away in dark rooms has already been well answered by Friston et al. (2012a). On the whole embodied creatures expect to stay warm, well fed and healthy. Dark rooms are not the kinds of environments that allow living agents to meet these basic biological needs. An agent that “felt the pull of the dark room” (Clark 2016) would be an agent that would after a while experience dehydration and hypoglycaemia, bodily conditions that are highly surprising.

While this response is clearly correct it only takes us so far. What is missing still is an explanation of the adaptive importance of things such as curiosity, play, and the spirit of adventure. Why would an agent that aims to occupy only those sensory states that are expected ever engage in behaviours that lead to novel and surprising discoveries, as happens when we play and explore? More specifically, why would agents ever be motivated to engage in such behaviours? If positively valued states are sensory states that are expected while negatively valued states are surprising, an agent whose actions are the result of active inference should only act to bring about sensory states that are expected. Novel sensory states such as those that occur as a result of exploration should be negatively valued (i.e. highly aversive). Yet this is not the case: many valued experiences are discovered by us through exploration.

2 The exploit/explore dilemma

In recent work, Friston and colleagues have shown that in active inference agents don't only act to keep themselves in states that are expected, but also act so as to minimise uncertainty about future outcomes (Friston et al. 2014, 2015, 2017; Schwartenbeck et al. 2013). They show how curiosity and novelty-seeking, and exploratory behaviour more generally, allow agents to reduce or resolve uncertainty about the world. Consider a scenario in which the agent is uncertain about which outcome to prefer, such as a mouse in a maze that needs to find its way to an unknown location of a reward while avoiding harm along the way. Recall that in active inference, preferences take the form of sensory states the agent expects to occupy or regularly frequent over time. The mouse is uncertain about where the dangers lie, and where the rewards are to be found. Its priors therefore tell it to keep its options open. It should resolve its uncertainty by further exploring the maze.

Friston and colleagues call priors that inform agents about which states they can expect to frequent over the long run “policies”. An “action policy” as Friston and colleagues use this term, can be thought of as a rule for the selection of a sequence of actions. We could think of policies in terms of paths of activity some of which are more probable than others to lead you to be in the states you expect to be in. Policies

should serve to minimise future free energy, or what Friston and colleagues understand in terms of “expected free energy” (Friston et al. 2015, 2017). Expected free energy is the free energy an agent expects to receive for each of its different policies, were it to pursue them (i.e. the trajectory or sequence of sensory states it expects in the future as a consequence of its actions).⁵ The only prior belief about a policy that is consistent with an agent’s continued existence is that it will pursue policies that reduce expected free energy. An agent that didn’t select policies based on such a prior would be unable to stay well adapted to its niche, and would eventually cease to exist.⁶ This implies that policies that have the highest prior probability are rules for generating action that will most likely help the agent to attain what they want in the future.⁷ This is because free energy is the divergence between the states an agent predicts it is likely to occupy (its posterior predictions) and the state it believes it should occupy if it is to satisfy its preferences (i.e. the states it expects to be in over time). The higher the prior probability of a policy, the less free energy an agent can expect in the future. When a policy has a high prior probability, the agent can allow the policy to drive action and be confident it will attain what it expects. We see this, for instance in the case of habitual behaviours. These are behaviours we have performed on many occasions that reliably lead to comfortable and familiar outcomes. They are policies that for this reason have a high probability because the free energy we can expect from following them is low.

According to Friston and colleagues, FEP mandates that an agent’s choices should reflect their beliefs about which policies have the highest prior probability of causing them to occupy the states they expect in the future. These beliefs inform the agent about the probability of reaching the states it expects from the states it currently occupies. An action policy will be assigned a value based on how well it is predicted to do at minimising the divergence between the states an agent is likely to occupy, and the state it believes it should occupy (i.e. its desired states such as securing a maximum payoff in a game). When the pursuit of an action possibility does not lead to the consequences an agent expects, she will engage in an epistemic action and explore her environment. By contrast, when the pursuit of action possibilities does lead to the consequences an agent expects, and a policy allows an agent to predict a clear path from her current states to the states she wants to occupy, she will treat the policy as highly probable. Policies that the agent believes have a high probability (such as habits) drive an agent’s actions since they stand the best chance of minimising free energy.

Friston and colleagues characterise the agent’s beliefs about the probability of a given policy in terms of the “precision” of a policy (Schwartenbeck et al. 2014). Precision weighting adds a second-order layer to active inference. In addition to esti-

⁵ Technically, the expected free energy also includes an ambiguity term; in other words, the expected free energy or uncertainty comprises the pragmatic value of behaviour in relation to prior preferences plus an epistemic term that resolves ambiguity about the causes of sensations. We thank an anonymous reviewer for drawing our attention to this important detail.

⁶ Based on this assumption that the probability of a policy is proportional to expected free energy, Friston and colleagues have built simulations of economic decision making and foraging behaviour (Friston et al. 2014, 2015, 2017).

⁷ It should be noted that decision-making is always playing out in a hierarchically organised biological architecture in which there are multiple policies in play operating over multiple time scales. This reflects the individual’s simultaneous openness to multiple relevant affordances (Bruineberg and Rietveld 2014).

imating the probability distributions of outcomes the predictive brain must also track the reliability (or precision) of its own estimates given the state of the organism and the current context. It uses this estimation of reliability to flexibly adjust the gain (or the “volume”) of particular error units: increasing the impact those units will have on the unfolding process (Friston 2010).⁸

The agent’s overall confidence in a policy is reflected in the *precision* of its policies. This precision is updated as a function of the predictive success of the consequences of actions. When the consequences of acting on a policy are correctly predicted (i.e. the potential outcomes are clearly consistent with prior preferences), precision increases and it decreases when new sensory input is surprising. In other words, if things are unfolding as expected, we become increasingly confident in the policy that we are pursuing. The precision of our beliefs about our own behaviour will increase as expected free energy decreases.

Low precision skews the prior over policies to a flat distribution, which reflects the likelihood that policies will be explored with roughly equal probability. Low precision causes the agent to keep its options open, and explore different options (as in the maze example sketched above). High precision by contrast, skews the prior to those policies that have the lowest expected free energy.⁹ It is the precision of beliefs about competing policies that Friston and colleagues hypothesise will decide whether an agent continues to follow a well-trodden path or departs from this path to actively explore the world.¹⁰ We can think of epistemic actions as being selected based on a recognition of the current state of the world as offering what one might call epistemic possibilities for action or “epistemic affordances”. The world is recognised as offering epistemic affordances when to put it in Friston’s terms (1) there is uncertainty to be resolved and (2) there is a clear and precise way forward that is driven by our beliefs about the policies that will best minimise expected free energy.¹¹ An example might be the uncertainty one experiences about the colour of an item of clothing due to artificial shop lighting. One might resolve this uncertainty by say asking the shop assistant if one can take the item of clothing out of the shop to view it under natural light, or by comparing it with other items whose colour one is more certain about.

⁸ Much of the argument in this paper is based upon formulations of active inference in terms of expected free energy using discrete state space (Markov decision) processes (MDP). The concept of prediction error units appeals to a slightly different formalism; namely, predictive coding and Bayesian filtering. However, it is possible to formulate the belief propagation in discrete (MDP) schemes in terms of prediction errors. We can therefore borrow the notion of attention and precision as it is used in predictive coding, when referring to active inference in choosing discrete policies, with no loss of generality. We thank an anonymous reviewer for drawing our attention to this difference.

⁹ It should be noted that epistemic actions are not always associated with low precision. If there is a clear and obvious uncertainty-reducing sequence of actions available, precision will increase to effectively render that epistemic course of action more likely to be selected. We thank an anonymous reviewer for emphasising this point to us.

¹⁰ We thank Jelle Bruineberg for clarifying discussion of Friston and colleagues’ complex work on epistemic value and for helping us fine tune some of our formulations of these ideas in the text.

¹¹ We thank an anonymous reviewer for this suggested characterisation of epistemic action in terms of recognition of epistemic affordances. For previous discussion of the concept of epistemic affordances, though not in the terms of Markov decision processes we are employing in this paper (see Rietveld & Kiverstein 2014).

While Friston and colleagues have provided an elegant set of formal tools for explaining exploratory behaviour and how agents resolve the so-called “explore-exploit” dilemma, their proposal nevertheless leaves us with two questions unanswered. It is these questions we take up in the remainder of our paper.

First, play and curiosity result in surprising and unexpected discoveries which are often experienced as having positive hedonic value. The novel experiences we have as a result of exploring our environment often feel pleasurable. Think of visiting a new culture as an example. The positive feelings of pleasure are part of what motivate us to engage in this type of exploratory activity rather than sticking with the comfortable familiarity of what is already known. Friston and colleagues suggest that this feeling of pleasure should be a consequence of recognising the action opportunities the world offers to reduce or resolve uncertainty (i.e. an epistemic affordance of the agent’s situation). Recall how the states an agent values are unsurprising states the agent expects to occupy. These states are the consequences of the behaviours the agent performs in seeking to minimise expected free energy. In other words the states an agent expects to occupy given its priors should be states with positive hedonic value. This however leaves it unexplained why novel experiences that reduce uncertainty should *feel good*. Why should there be a positive phenomenology that comes with exploring and making progress in reducing uncertainty? We enjoy being curious. This is part of the value we assign to these activities. Insofar as current work on epistemic action doesn’t account for the positive feelings that characterises curiosity and play, it doesn’t yet fully account for how value works in epistemic action.

Second, the success of Friston and colleague’s account depends on their explaining how it is that agents are able to optimise the precision of their policies. They say that precision estimation is a consequence of free energy minimisation. They show how the optimisation of precision estimations has dynamical properties that closely resemble those of dopaminergic systems in the brain.¹² This provides a candidate mechanism for the implementation of precision in active-inference. However it leaves us with questions still about how precision is weighted in a given context. We will show how these two questions may turn out to share a common answer. Bodily feelings in the form of affordance-related states of action readiness turn out to be a part of what allows an agent to estimate precision on the fly.

3 The feeling of action readiness

We’ve seen in the previous section how expected free energy gets to decide whether the agent exploits familiar solutions or explores the environment to reduce uncertainty. Precision influences how evidence is accumulated and thus how “beliefs” are formed by an agent. Precision estimates are made more generally by the brain to separate out organism-important error signals from the surrounding unimportant noise (Feldman 2013). For example, while walking home on a familiar busy street one might barely

¹² It has been suggested that precision weighting is the result of neuromodulators (such a dopamine). Kanai and Friston have suggested that, “this may explain why superficial pyramidal cells have so many synaptic gain control mechanisms such as *N*-methyl-D-aspartate (NMDA) receptors and classical neuromodulatory receptors like D1 dopamine receptors” (Kanai et al. 2015).

notice the general buzz of the people around you. Day to day the buzz is inevitably different in shape—different people interacting in different ways. Nevertheless, while the exact shape of the buzz is exceptionally unpredictable it draws almost no attention: no further processing is allocated. Now if someone were to unexpectedly fall close by and seem to require assistance this bit of error would suddenly be pertinent. If we think of error signals as broadcasting newsworthy information then we must also explain how the brain decides which channels it should “listen to” (Kanai et al. 2015; c.f. Kwisthout et al. 2017). Estimating the precision of a policy is a special case of a more general phenomenon in which the agent is continuously engaged in monitoring its own level of confidence in its predictions about the world.

The foregoing concerns the precision of prediction errors in relation to states of the external environment. However we have been discussing precision in relation to action-selection, and have therefore been discussing the precision of beliefs about policies. These notions of precision turn out to be closely related and intertwined on our account.¹³ Policies consist of interrelated, and nested states of action readiness, which are patterns of readiness for the sensory consequences or outcomes of action. We therefore suggest that in assigning precision the brain isn’t *only* concerned with the reliability of the PE signal. It is more accurate to say that precision relates to how well the agent is doing at engaging with expected uncertainty in relation to the sensory consequences of temporally extended sequences of actions. Precision doesn’t just concern the agent here and now and its momentary state of uncertainty with regards to some current prediction error. We’ve seen above how expected or future uncertainty is also important when it comes to assigning precision to policies. FEM agents actively seek a means of managing uncertainty over time. The rollercoaster of continual increases and decreases of errors that accompany life become expected and are folded into our expectations. For such a system it becomes important not only to track the constantly fluctuating instantaneous errors, but also to pay attention to the dynamics of error reduction over longer time scales. This means paying attention to *the rate* at which those errors are being reduced or increasing. Rate of change is an important (but largely overlooked dimension) of free energy minimisation. If we compare two agents both of which succeed in dealing with prediction error the agent that does it faster will do better in the long run than the agent that takes longer.¹⁴

We can think of the rate of change of prediction error reduction by analogy with velocity (Joffily and Coricelli 2013: p. 3). The velocity of an object is the rate of change in the position of an object relative to a frame of reference over time. So velocity is equivalent to the speed of an object moving in a particular direction. Rate of change in relation to prediction error reduction thus refers to how fast or slow prediction error is being reduced relative to the states of the whole agent-environment system. If the speed of error reduction increases, this equates to decrease in free energy over time

¹³ See our earlier footnote 8.

¹⁴ Technically, what we are saying here is that the path or time integral of prediction error (or expected free energy) is the important thing. Crucially, because these time integrals or averages are known as Hamiltonian Action, we are saying is that the best policies will conform to Hamilton’s Principle of Least Action. We thank an anonymous reviewer for the formulation of this technical point, and for making the connection to Hamilton’s principle.

(relative to what was expected). If speed of error reduction decreases, this equates to an increase in free energy over time.

Each agent's performance in reducing error can be plotted as a slope that depicts the speed at which errors are being accommodated relative to their expectations. The steepness of the slope indicates that error is being reduced over a shorter period of time and so faster than the agent expected: the steeper the slope, the *faster* the rate of reduction. Think of mastering a second language and finding it easy to take part in a conversation with a stranger in this second language. A gentle slope by contrast indicates that error is being reduced at a *slower* than expected rate. The agent has encountered an error that is proving difficult to deal with. This has the result that they reduce fewer errors over time. Suppose for instance that a person has broken a leg and they now need to get around their environment on crutches. They are now much slower to move around—the rate at which they are able to get into the states they expected to be in has slowed dramatically. This is typically the source of some frustration for the person.

Once we have the idea of rate of change in play, we see that even large instantaneous errors (sudden spikes in the slope) could be experienced as positive as long as the error takes place within a more general reduction of error over time. This is to say that a large but resolvable error signal informs the agent that the environment offers the opportunity to resolve uncertainty. An epistemic action then becomes an attractive option.

This goes a long way to helping make sense of why certain errors are acceptable and even highly desirable. The environment offers an epistemic affordance, an opportunity for information gain that allows one to pick up on a free-energy lowering policy. Information about rate of change is thus highly important for updating one's expectations, changing the course of action entirely or rather continuing on the same course of action. It informs the agent whether, given what they know already, there is still room to improve. Perhaps there is no room for improvement because error rate just keeps increasing in which case what you should do is just try something different. You shouldn't change what you anticipate happening when the result of making such a change will just be more error. You should instead explore and look for new opportunities that help you to learn to grip better.

In the recent literature a number of authors have begun to relate information about rate of change in error reduction to the valence of full-blooded emotional experience like happiness, disappointment, hope and fear (Joffily and Coricelli 2013; Van de Cruys 2017; Van de Cruys and Wagemans 2011). These authors hypothesise that the valence of different emotional experiences is a reflection of *unexpected rate of change* in prediction error reduction. When free energy is increasing at a greater than expected rate this can feel bad, it can for instance be experienced as frustrating. Conversely when free energy is decreasing at a faster rate than expected this can feel really good. For example, a positive emotion like happiness is a reflection of an unexpected reduction of prediction error (e.g. error being reduced at a rate faster than expected), while a negative emotion like disappointment reflects an unexpected decrease in prediction error reduction (e.g. error being reduced at a rate slower than expected).

In this section we've been applying the notion of rate of change to explain why an agent might be moved to explore rather than exploit. Assuming these authors are

right to tie valence to rate of change, our proposal is that it is valence that plays this role of motivating an agent to engage in exploratory actions rather than exploit.¹⁵ The concept of “valence” is used by these authors to refer to the positive or negative (approach or avoid) character of an emotional experience that inform the organism about its current relationship with the environment. Pleasurable states have positive hedonic value and are associated with approach behaviours. Aversive states have negative hedonic value and are in turn associated with avoidance behaviours. Their claim is not so simple as to say we feel good when our predictions fit and bad when they do not. They claim (rightly) that the predictive organism is in a constant state of error management at all levels of the hierarchy, and yet clearly there is not a felt experience of each and every fluctuation. Errors are reviewed on a background of learned expectations concerning how fast or slow (the rate) such errors have been reduced previously (Joffily and Coricelli 2013).

We suggest thinking of valence differently in terms of multiple states of affordance-related action readiness that are simultaneously affective and behavioural (Rietveld 2008).¹⁶ At the same time as emotional experiences feel good or bad they also prepare or make us ready to act on relevant affordances (possibilities for action offered by the environment). It is the relevant affordances of the environment that have valence. This valence consists in solicitations or invitations to act: relevant affordances attract or repel the agent’s actions (Bruineberg and Rietveld 2014). Think of how the apple sitting next to you as you work can look enticing to eat when you are hungry: it has positive valence. But when you take a bite into and find it is rotten inside it ceases to be enticing in the same way and takes on a negative valence. We thus disagree with Joffily and Coricelli (2013) and Van de Cruys (2017) who treat valence as the avoid/approach character of full-blooded emotional experiences. They treat valence as a property of emotional experience while we suggest understanding valence in relational terms, and as necessarily environment involving: it is relevant affordances that have valence in virtue of which they solicit or invite some form of action on the part of the agent.

¹⁵ A closely related proposal can be found in Van de Cruys (2017). He emphasises that valence as error dynamics helps guide the organism in its exploration to niches that have maximum gain in prediction-error minimisation. Van de Cruys doesn’t explicitly apply this hypothesis to the case of epistemic action as we have been doing in this paper. However, we can suppose that sometimes increasing PE may lead an agent to shift into an exploratory mode of engagement with their environment, while on other occasions it may lead the agent to shift to another reliable exploit mode. Exactly what the agent does next will depend on the precision of the policies it is enacting as discussed above. Differences between our ecological-enactive account of valence and the view that can be found in Van de Cruys are discussed below.

¹⁶ Here we are drawing upon the analysis of emotion found in the research of the emotion psychologist, Nico Frijda (see Frijda 1986, 2007), in which states of action readiness are central. We will have nothing to say here about so-called *full-blooded* emotions that Frijda analyses (e.g. happiness, sadness, anger, etc.). Our interest will be instead in the myriad ways that the brain and body change as part of our meaningful engagement with the world. Our use of felt states of action readiness to refer to the individual agent’s openness to a field of relevant affordances closely parallels Matthew Ratcliffe’s account of “existential feelings” (2008). Existential feelings are background bodily states of action readiness that orient us in the world. They are what attune us towards or away from relevance. While existential feelings are certainly felt, the affective changes can themselves be more or less explicit, foregrounded objects of attention. So for example, we may be expressly aware of our experienced anxiety, or we may simply perceive a situation (through our anxious state) as worrisome.

In addition to valence at the level of individual relevant affordances we suggest it also makes sense to think in terms of the field of relevant affordances as a whole as having valence. This is the kind of valence that is made explicit when someone asks how things are going, or how one feels in a situation. One's initial response to this question relates to the situation as a whole, though of course it is possible to zoom in on particular aspects to make things more specific. The field of relevant affordances comprises the multiple relevant affordances that get the agent bodily ready to respond. One is ready to respond to each of the relevant affordances but one also has an overall grip on the situation. It is this valence at the global level of the field as a whole that we propose to understand in terms of rate of change. Global level valence gets the agent ready to either exploit the particular inviting relevant affordances or to seek out epistemic affordances that offer opportunities for information gain, thereby allowing one to pick up on a free-energy lowering policy.

Valence can thus be thought of both at the global level of the agent in relation to the field as a whole, and at the more local level of micro-level states of action readiness elicited by relevant affordances that invite the agent to act. The former ("global valence") is best thought of in the context of what we've earlier described as the tendency towards an optimal grip on the field of relevant affordances (Bruineberg and Rietveld 2014). The relation between multiple micro-level states of action-readiness and macroscopic patterns of activity at the level of the individual agent as a whole is analysed in terms of self-organising dynamics in our earlier work (Bruineberg et al. 2016; Rietveld, Denys & van Westen, 2017).

In earlier work (Bruineberg and Rietveld 2014; Bruineberg et al. 2016), we've proposed an ecological-enactive reading of the FEP, providing an analysis of free-energy in terms of disattunement of internal and external dynamics, a dynamic state of disequilibrium within the agent-environment system as a whole. In active inference, agents prepare actions that will reduce disattunement and thereby lead them closer towards some dynamical equilibrium or grip on the situation. States of action readiness originate in fluctuations of affect that orient us towards affordances that matter to us, preparing us for sensory consequences that arise from responding to inviting possibilities for action. Emerging a few hundred milliseconds after an event, action readiness is the earliest coordinated evaluation of the new situation by the organism as a whole (Klaasen et al. 2010). Positive and negative feeling should be thought of as an integral part of tending towards an optimal grip. As Frijda notes:

"Emotional feeling is to a very large extent awareness, not of the body, but of the body striving, and not merely of the body striving, but the body striving in the world... emotional experience is to a large extent experienced action tendency, or experienced state of action readiness." (Frijda 2004: p. 161, quoted by Lowe and Ziemke 2011: p. 8)

Felt states of action readiness manifest as a "complex space of polarities and combinations" (Colombetti 2005; Thompson 2007: p. 378). The agent prepares to move towards/away, approach/withdraw, and is receptive/defensive to affordances in the environment that are exerting a pull on, or repelling the agent. These movement tendencies can also be consciously felt as pleasant/unpleasant, positive/negative, and they relate to affordances the individual likes/dislikes, is attracted/repelled by. We suggest that felt states of action readiness arise when there is an unexpected change in rate of error reduction at the level of the agent-environment system as a whole. Rate of

change is an important source of information for the agent because it can help them to always be ready for opportunities for improving grip, and living systems continuously strive to improve grip (Bruineberg and Rietveld 2014).

This ecological-enactive analysis of rate of change can help us to address the first of the questions we raised about epistemic action in the previous section. There we raised the question of why reducing uncertainty through epistemic actions should feel good to the agent. Why should novel experiences have a positive phenomenology? We've just suggested thinking of rate of change as the changes felt in the skilled body (in its relation to the field of relevant affordances as a whole) with a positive or negative hedonic value. When an agent succeeds in reducing error at a faster than expected rate (or recognises the opportunity to do so) this feels good. It is thus not uncertainty reduction alone that the agent cares about, but also the rate at which uncertainty is being reduced. Pleasurable feelings arise in the form of feedback as part of the process of our moving towards, or being drawn towards affordances that are relevant to us. We are drawn towards opportunities for improving grip, and positive feelings arise when we improve grip at a faster than expected rate. This is to say that the slope one could plot describing the rate of error reduction would have a steep incline.¹⁷

Conversely, negative affect is experienced by the agent in terms of being repelled from a situation. This occurs when we do worse (or anticipate doing worse) than expected at reducing error, and the slope describing error reduction has a gentle incline. Consider boredom as an example. Suppose you find yourself stuck in a seat in a music hall sitting through a boring symphony. The music you are hearing has nothing to offer. It is boring because there are no salient regularities to be harvested relative to the skills you have that attune you to the environment. There is no opportunity to do better than expected at reducing prediction error because the skills you have already attune you to the music without this requiring any effort from you. Alternatively the skills of the individual may be such that the structure of the music is too complex to get a grip on it. This is an inherently frustrating situation for the listener stuck in their seat who nevertheless aims to continuously improve their grip on the world. The felt frustration manifests as boredom for the music and agitation, a drive to get up out of your seat and leave.

In the next section we argue that rate of change may also play a key role in *precision estimation* as it occurs in active inference. Active inference can be tuned by rate of change in agents that are sensitive to this feedback signal. They can use this feedback signal to allow themselves to be pushed towards or pulled away from aspects of the environment that offer opportunities for making progress in uncertainty reduction. The hypothesised role of rate of change in precision estimation is among the important novel contributions of this paper.¹⁸

¹⁷ Again we see Hamilton's Principle of Least Action in play. Our thanks to an anonymous reviewer for this point.

¹⁸ It is a possibility that has not to our knowledge been recognised in previous work that has appealed to error dynamics to explain the valence of emotional experience.

4 Using rate of change to tune precision-weighting on the fly

How is precision weighted in active inference? We suggest in line with our ecological-enactive reading of FEP that precision should be understood in the context of tending towards an optimal grip on the affordances available in the ecological niche. Precision is what sets the degree of influence on behaviour (or more precisely, action readiness) of the multiple relevant affordances inviting us to act.¹⁹ Error dynamics (rates of FEM) are “grasped corporeally”, we feel the dynamics of error-reduction (and whether it is going well or badly) in our attunement to the world (Patočka 1998). The agent doesn’t simply act so as to improve grip; it is a part of their acting skillfully that they can do so with sensitivity to how well or badly they are doing. Agents that make use of feelings that arise from rates of change can continuously do better at improving their grip on what is relevant in the landscape of affordance by exploring and seeking out novelty. They can aim to better engage with error, and attune to the unexpected so as to broaden their skills and grip in more and more domains of their ecological niche.

Once we understand FEM as always being enacted in an ecological context as we propose, it makes sense to suppose that in general, FEM agents will be on the lookout for opportunities that are rich in the kinds of error they can manage given their skills (eg. manageable errors). We suggest this is the kind of error that will allow an agent to improve on their level of skilled engagement with affordances. Unexpected improvements or setbacks, in relation to the expected rate of error reduction, then provide a particularly valuable learning signal that can direct resources to opportunities for improvement or speak in favour of task-switching (see Rietveld and Brouwers 2016 for real life examples of this). FEM organisms do *not* try to *maximally* reduce error, since sometimes error can be invaluable for learning,²⁰ and in any case prediction error is unavoidable. When we think about a day that has gone well, we evaluate the day as a whole in part based on the surprising and unexpected things that happened to us and how well we managed to deal with them. A good day is not only one in which we succeeded in reducing prediction error relative to our expectations. It is one in which we were met with all manner of unexpected events and we did well at meeting these challenges.

Agents can be sensitive to different degrees to how well or badly they are doing at gripping to their environment. To put this in terms of rate of change of FEM, they are sensitive to different degrees to the increases and decreases in FEM. Given an advanced level of skill we expect our skilled body (as a model of the world) to do well at attuning to a given context. When we run into more troubles than we anticipated, this slows down the rate at which our current skills succeed in attuning us to unexpected changes in the environment. Think about a difficult day at work in which many problems come up that you fail to solve. The feeling of frustration in this case is feedback that informs us that things are going worse than expected. This is important information for an agent because it can be used to move us to do things differently. When free energy is rapidly increasing, this is a sign that the agent is doing poorly over time

¹⁹ This hypothesis about the precision mechanism is developed in more detail in Miller and Clark (2017).

²⁰ Schouppe et al. (2014) shows that more reward comes from resolving incongruent stimuli than from congruent stimuli in a Stroop task. These findings points to initial errors being conducive to higher rewards.

at accommodating sensory input. An agent that possesses this information about rate of change should downgrade confidence in their policies. Equivalently, they should assign more confidence to prediction errors relative to their policies.

Now consider a more positive scenario in which prediction error progressively decreases like in the example of the day in which lots of unexpected events occur that we nevertheless manage well. The agent's policies are performing well at attuning them to change in the environment, and thus they have good cause for being highly confident in their policies.

The more an agent takes note of rate of change, the more sensitive they are to how well they are gripping in a given situation. Agents that lack sensitivity to felt states of action readiness will be more likely to get stuck in situations that are frustrating to them. Feelings provide them with the impetus or impulse to switch. When people are not sensitive to these feelings, this can lead them to be overconfident or underconfident in their policies. Overconfidence in one's policies can lead one to overlook unexpected changes in the environment that bodily feelings attune us to. Underconfidence in one's policies can make an agent dissatisfied with what they are doing when they are doing just fine at adapting to the unexpected. Sensitivity to rate of change and to the feelings it gives rise to can be used to tune confidence as we go.²¹ This is maybe one of the essential ways in which agents are able to stay attuned to what matters to them. Failure to tune confidence as we go along leads to inflexibility and failure to switch activities based on changes in the environment or changes in internal state (e.g. homeostatic needs). An agent might for instance persist with some activity that has been weighted as highly important, but is failing to reduce free energy at the expected rate because of some other pressing need that is being neglected. Sensitivity to rate of change, which reflects how things are going at the global level of the agent as a whole, would tell the agent that they need to assign priorities differently allowing them to do better at reducing free energy overall by continually tuning their confidence in the expectations that are driving their actions on the fly.²²

Why are people so curious and playful? We can remove some of this mystery we suggest, once we appreciate the role of rate of change in precision estimation. Agents that weight precision based on feedback from the feeling of grip will be attracted to opportunities for continually improving in their skills. Positive and negative hedonic value is felt in the body and can (in the agent that is sensitively attuned to these

²¹ Although beyond the remit of our current treatment, it is a prediction of our arguments that the dopamine system should play a central role in this online tuning of precision in active inference (thanks once again to an anonymous reviewer for suggesting this point to us). This is because, in physiological versions of predictive coding and active inference, precision is thought to be mediated by neuromodulatory effects (see e.g. Friston et al. 2012a). Key among these (in the domain of motivated behaviour) is dopamine—a neurotransmitter strongly implicated in hedonics, reward and emotional learning.

²² Inflexibility that arises from failure to tune confidence on the fly may go some way towards accounting for the current political climate in which people are falling back on comfortable, familiar certainties from the past in response to the many complex, hard to manage uncertainties of the future. The current political situation presents people with a tremendous amount of uncertainty that is far too complex for them to integrate. One natural response to this is to fall back on outdated but familiar and entrenched patterns of interaction that no longer give a grip on current events, which for interesting sociological reasons people nevertheless overlook.

feelings) provide feedback that moves the agents through the environment leading them to places where they stand to learn the most.

5 Novelty seeking and learning progress

We've been arguing it is not only prediction error that the agent seeks to reduce in their skilled engagement with the environment, but also the rate of change in prediction error (i.e. the opportunity to reduce uncertainty in the future). Anticipation isn't just about determining what is most likely to happen next. Optimising one's engagement with a dynamically changing environment requires in addition, sensitivity to how well one is doing at reducing disattunement over time. Sometimes it feels good for the agent to generate more prediction error in their interaction with the environment as a part of their epistemic foraging.

Irreducible error means environmental complexity is too high for the agent—and this feels bad, as in the example of the symphony that one finds boring because its complexity is too high. Too little error in our dealings with the environment means that our model is already fitting well to the environment, there is nothing further to be learned and we feel bored. Value in epistemic action thus seems to be a matter of finding the right balance so that the agent is continuously improving in the speed at which they are reducing prediction error.

Imagine signing up for a one hour swimming lesson to improve one's swimming skills. When one arrives for the first time at the swimming pool for the training one finds out that the group is very large and that it is not a beginners group. Most people in the class are already proficient at doing breast-stroke for instance, something one has not yet mastered. Given the size of the class, the teacher probably notices that one's level of performance is lagging behind other members of the class but is unable to give one the attention needed for acquiring a new skill. One feels that things are going a good deal worse than one expected, and one's swimming ability is not at all improving. Every new exercise during the class is too difficult. The valence of this situation is negative; one feels out of place and has the action tendency of leaving the class, and joining a different one that is better suited to one's level of ability. Rather than being determined to keep trying until one has mastered the exercises sufficiently, one might well decide it would be better to switch and find a class that is targeted at beginners and has more personal attention for the participants.

The richest opportunities for improving in FEM will come from situations that are neither too complex, nor so simple and straightforward that we already know how to deal with them. In finding this balance of complexity and simplicity an agent is able to learn optimally and so is able to do better, while nevertheless always falling short of fully attaining an optimal grip or equilibrium with the environment. This means being sensitive to the felt states of action readiness that attune one to positively or negatively valenced relevant affordances. Negatively charged feelings of action readiness tell one that things are not going well. One is failing to grip in the way one expects given one's level of confidence. This may make switching activities or strategies an enticing option so long as one is sensitive to this feeling. Positive feelings tell you that you are

doing well at dealing with, and adapting to unexpected changes. This provides you with valuable feedback for further improving your skills.

Various lines of research support the view that agents seek out environments with optimal amounts of novel complexity (error). For example, [Berlyne \(1966\)](#) argues that organisms actively seek out stimuli that are slightly above the complexity the organism is used to. A few good examples of this come from research on early childhood development in which newborns were found to attend longer to stimuli that are neither too simple nor too complex ([Kidd et al. 2012](#)). In non-human studies, rats were found to frequent parts of a maze that were decorated in a slightly more complex fashion than parts they commonly frequented in the past ([Dember et al. 1957](#)). Learning environments in which the complexity is just above the abilities of the agent offer the largest accelerations of error reduction. These are the environments the rats *like* to explore the most because they offer the most opportunities to learn relative to what they know already. Exploration of these environments feels good to the rats because it is in such environments that acceleration of FEM is at its greatest. We hypothesise that the rats are motivated to explore because it feels good to explore this kind of environment—doing so offers them the greatest opportunity for acceleration of FEM.

A similar line of thinking comes from recent research on error reduction dynamics in artificial intelligence and robotics ([Oudeyer et al. 2007, 2013](#); [Schmidhuber 2010](#)). Kaplan and Oudeyer take the rate of error reduction to be associated with intrinsic rewards in humans ([2007](#)).²³ By linking error reduction dynamics and intrinsic rewards they offer a model of learning in which agents are intrinsically driven to investigate particular regions of the environment as long as there are learnable regularities left to harvest given their current skill level. Being sensitive to error dynamics guarantees that the agent avoids wasting time in places where regularities are either already learned or too complex given the agent's skill level. To put this in the terms of our paper, sensitivity to felt states of action readiness tunes the agent to, and draws them to explore, learning-rich places simply by tracking local learning progress. Such systems will naturally and spontaneously move from one stage of development to the next, from one level of complexity to another, as error reduction becomes less available (either all that's left is uninteresting noise or the complexity is yet too high to be managed given the skill level of the organism). A neat outcome of such systems is the self-organization of developmental and learning trajectories that naturally move agents from acquiring simple to more complex skills over time ([Oudeyer and Kaplan 2006](#); [Oudeyer et al. 2007](#); [Kaplan and Oudeyer 2011](#); [Moulin-Frier and Oudeyer 2012](#)).

These ideas about learning progress have shown promise in developmental robotics, allowing robots implementing these routines to “efficiently learn repertoires of skills in high dimensions and under strong time constraints and to avoid unfruitful activities that are either well learnt and trivial, or which are random and unlearnable ([Pape et al. 2012](#); [Ngo et al. 2012](#); [Baranes and Oudeyer 2013](#); [Nguyen and Oudeyer 2013](#))” ([Gottlieb et al. 2013](#): p. 9). There are good reasons to suspect a similar approach to

²³ In the free energy framework, the concept of intrinsic motivation that underlies (self) exploration (e.g. motor babbling) and novelty seeking is understood in terms of *expected free energy*; also referred to as Bayesian surprise, information gain in the literature. Intrinsic value is the value of information or epistemic value (see Sect. 2 above).

learning takes place in humans. Progress-based systems in robotics simulate closely infant sensorimotor development. This has led to the hypothesis that certain patterns of information-seeking behaviour in humans may emerge from a particular embodied system (morphology, etc.) intrinsically motivated by progress-based learning strategies (Smith 2003; Kaplan and Oudeyer 2007; Oudeyer et al. 2007). To put this point in our own words: the sensitivity to the valence of a situation moves agents to explore situations that maximize learning progress by directing them towards a trajectory of action possibilities whose complexity is neither too simple nor too novel. Oudeyer and colleagues refer to these learning sweet-spots as *progress niches*:

“Progress niches are not intrinsic properties of the environment. They result from a relation between a particular environment, a particular embodiment [...] (sensors, actuators, feature detectors, and techniques used by the prediction algorithms), and a particular time in the developmental history of the agent. Once discovered, progress niches progressively disappear as they become more predictable.” (Oudeyer et al. 2007, p. 282)

By tracking their own progress in learning, agents are moved to seek out opportunities that present them with just the right level of complexity that they can make something of these opportunities given their current level of ability.²⁴ FEM agents don’t need to orient towards novelty indiscriminately. Agents that make use of their sensitivity to how well they are doing at tending towards an optimal grip can orient to *the right kind of novelty*—the kind that maximises their own learning rate or the opportunity to do better at improving their grip on what matters in the environment. They will be agents that are intrinsically motivated to explore, seek out novelty and along the way improve their skill at gripping.

To offer a rather extreme example of this, some people are able to forecast events like election results, economic collapses, famines and wars with an accuracy much better than chance. They have been shown to perform on average 65% better than the average person and 30% better than US intelligence agents that forecast these kinds of events for a living (Tetlock and Gardner 2015).²⁵ These people are highly skilled at estimating the probable outcomes of counterfactual scenarios. It is a prediction of our arguments that the secret to their success must lie in their ability to form better expectations than the average person about how errors will arise. They do a way better job than the average person of anticipating when the unexpected is likely to arise and adapting their probability estimates accordingly. If our arguments are along the right lines, this is something they are able to do by paying close attention to what their feelings tell them to do so they can continuously tune the confidence they have in their own estimations of what could happen in the future. They are always looking for opportunities to incrementally improve in how they are doing at reducing error, and they don’t care so much about the end product. They are as Andy Clark nicely put it to us in conversation “slope-chasers”.

²⁴ We don’t mean to suggest that progress niches are discovered by individual agents all by themselves. Often other people from the practice, more experienced practitioners, give the most relevant feedback about how well one is performing and educate one’s attention to relevant affordances.

²⁵ Many thanks to Andy Clark for suggesting this fascinating example of super-forecasters.

Conclusion

We started the paper by raising a puzzle for FEP about why an agent would be motivated to play and explore if all they ever aim to do is keep themselves in states that are expected relative to its model of the world. We've argued that a FEM agent should naturally engage in epistemic foraging, and seek out epistemic affordances or opportunities to reduce uncertainty in relation to their policies. We've raised two questions for an account of epistemic foraging in terms of FEM and active inference. The first question asks why it should feel good to an agent to engage in exploratory, curiosity-driven behaviours. We've offered a new perspective on recent work on the rate of change in error reduction to address this question. According to this earlier work, it feels good for an agent to increase in the speed of error-reduction, and it feels bad for an agent when they reduce error at a slower than expected rate. We've suggested that what rate of change tracks are affordance-related changes in states of action readiness. Our contribution is thus two-fold. First we have proposed an ecological and enactive interpretation of these relatively recent ideas about what rate of change might be doing in predictive-processing. Second we have put this interpretation to work to explain the positive and negative hedonic value that can motivate an agent to engage in epistemic actions.

The second question we've raised asks how precision-weighting might work in active inference. Precision-weighting plays a crucial role in epistemic action since it is the precision that is assigned to a policy that settles the question for the agent as to whether to exploit what one already knows, or to explore and seek out novelty. Building on our ecological-enactive interpretation of FEM we've proposed that sensitivity to rate of error reduction may play a role in precision estimation. Active inference can be tuned on the fly in agents that use rate of change as a feedback signal, thereby guaranteeing that they continuously make progress in reducing their own uncertainty.

An organism wired to be rewarded for reducing disattunement at a faster rate should naturally and spontaneously orient itself to places in the ecological niche that offer opportunities where disattunement can be managed the best. They should be motivated to seek out opportunities for managing errors that are just above their current level of ability and to develop new skills. This is to say that the FEM agent ought to be a curious agent, motivated to explore and play in their environment, constantly pushing the boundaries of what they know how to do. In their exploratory engagement with their ecological niche, curious agents discover novel affordances that allow them to constantly improve in their skills.

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

References

- Baranes, A., & Oudeyer, P.-Y. (2013). Active learning of inverse models with intrinsically motivated goal exploration in robots. *Robotics and Autonomous Systems*, 61(1), 49–73.
- Berlyne, D. E. (1966). Curiosity and exploration. *Science*, 153(3731), 25–33. Jul 1.
- Bruineberg, J., & Rietveld, E. (2014). Self-organization, free energy minimization, and optimal grip on a field of affordances. *Frontiers in Human Neuroscience*, 8, 599.
- Bruineberg, J., Kiverstein, J., & Rietveld, E. (2016). *The anticipating brain is not a scientist: The free-energy principle from an ecological-enactive perspective* (pp. 1–28). Special Issue Predictive Brains: Synthese.
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *The Behavioral and Brain Sciences*, 36(03), 181–204. doi:10.1017/S0140525X12000477.
- Clark, A. (2016). Busting out: Predictive brains, embodied minds, and the puzzle of the evidentiary veil. *Nous*,. doi:10.1111/nous.12140.
- Colombetti, G. (2005). Appraising valence. *Journal of Consciousness Studies*, 12(8–9), 103–126.
- Dember, W. N., Earl, R. W., & Paradise, N. (1957). Response by rats to differential stimulus complexity. *Journal of Comparative and Physiological Psychology*, 50(5), 514–518. doi:10.1037/h0046780.
- Feldman, J. (2013). Tuning your priors to the world. *Topics in Cognitive Science*, 5(1), 13–34. doi:10.1111/tops.12003.
- Frijda, N. H. (1986). *The emotions: Studies in emotion and social interaction*. Cambridge: Cambridge University Press.
- Frijda, N. H. (2004). Emotions and action. In A. S. R. Manstead, N. Frijda, & A. Fischer (Eds.), *Feelings and emotions* (pp. 158–173). Cambridge: Cambridge University Press.
- Frijda, N. H. (2007). *The laws of emotion*. Mahwah: Lawrence Erlbaum Associates Publishers.
- Friston, K. J. (2009). The free-energy principle: A rough guide to the brain? *Trends in Cognitive Sciences*, 13(7), 293–301. doi:10.1016/j.tics.2009.04.005.
- Friston, K. J. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11(2), 127–138. doi:10.1038/nrn2787.
- Friston, K. (2013). Life as we know it. *Journal of the Royal Society Interface*, 10, 20130475. doi:10.1098/rsif.2013.0475.
- Friston, K. J., Fitzgerald, T., Rigoli, F., Schwartenbeck, P., & Pezzulo, G. (2017). Active inference: A process theory. *Neural Computation*, 29(1), 1–49.
- Friston, K. J., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., & Pezzulo, G. (2015). Active inference and epistemic value. *Cognitive Neuroscience*, 6(4), 187–214.
- Friston, K. J., & Stephan, K. E. (2007). Free-energy and the brain. *Synthese*, 159(3), 417–458.
- Friston, K. J., Schwartenbeck, P., Fitzgerald, T., Moutoussis, M., Behrens, T., & Dolan, R. J. (2014). The anatomy of choice: Dopamine and decision-making. *Philosophical Transactions of the Royal Society B (Biological Sciences)*,. doi:10.1098/rstb.2013.0481.
- Friston, K. J., Shiner, T., FitzGerald, T., Galea, J. M., Adams, R., Brown, H., et al. (2012a). Dopamine, affordance and active inference. *PLoS Computational Biology*, 8(1), e1002327. doi:10.1371/journal.pcbi.1002327.
- Friston, K. J., Thornton, C., & Clark, A. (2012b). Free-energy minimization and the dark-room problem. *Frontiers in Psychology*,. doi:10.3389/fpsyg.2012.00130.
- Gottlieb, J., Oudeyer, P. Y., Lopes, M., & Baranes, A. (2013). Information seeking, curiosity and attention: Computational and neural mechanisms. *Trends in Cognitive Science*, 17(11), 585–596.
- Hohwy, J. (2015). The neural organ explains the mind. In T. Metzinger & J. M. Windt (Eds.), *Open MIND: 19(T)*. Frankfurt am Main: MIND Group. doi:10.15502/9783958570016.
- Joffily, M., & Coricelli, G. (2013). Emotional valence and the free-energy principle. *PLoS Computational Biology*, 9(6), e1003094. doi:10.1371/journal.pcbi.1003094.
- Kanai, R., Komura, Y., Shipp, S., & Friston, K. (2015). Cerebral hierarchies: Predictive processing, precision and the pulvinar. *Philosophical Transactions of the Royal Society of London. Series B Biological Science*, 370(1668), 11–16.
- Kaplan, F., & Oudeyer, P. Y. (2007). In search of the neural circuits of intrinsic motivation. *Frontiers in Neuroscience*, 1(1), 225–236. doi:10.3389/neuro.01.1.1.017.2007.
- Kaplan, F., & Oudeyer, P.-Y. (2011). From hardware and software to kernels and envelopes: A concept shift for robotics, developmental psychology, and brain sciences. In J. L. Krichmar & H. Wagatsuma (Eds.), *Neuromorphic and brain-based robots* (pp. 217–250). Cambridge: Cambridge University Press.

- Kidd, C., Piantadosi, S. T., & Aslin, R. N. (2012). The Goldilocks effect: Human infants allocate attention to visual sequences that are neither too simple nor too complex. *PLoS ONE*, 7(5), e36399. doi:[10.1371/journal.pone.0036399](https://doi.org/10.1371/journal.pone.0036399).
- Kirchhoff, M., & Froese, T. (2017). Where there is life there is mind: In support of a strong life-mind continuity thesis. *Entropy*, 19(4), 169. doi:[10.3390/e19040169](https://doi.org/10.3390/e19040169).
- Klaasen, P., Rietveld, E., & Topal, J. (2010). Inviting complementary perspectives on situated normativity in everyday life. *Phenomenology and Cognitive Sciences*, 9(1), 53–73.
- Kwisthout, J., Bekkering, H., & Van Rooij, I. (2017). To be precise, the details don't matter. On predictive processing, precision and level of detail of predictions. *Brain and Cognition*, 112, 84–91.
- Lowe, R., & Ziemke, T. (2011). The feeling of action tendencies: On the emotional regulation of goal-directed behavior. *Frontiers in Psychology*, 2, 346. doi:[10.3389/fpsyg.2011.00346](https://doi.org/10.3389/fpsyg.2011.00346).
- McEvoy, P. (2002). *Classic theory: The theory of interacting systems*. San Francisco: Microanalytix.
- Miller, M., & Clark, A. (2017). *Happily entangled: Prediction, emotion and the embodied mind* (pp. 1–17). Special Issue, Predictive Brains: Synthese.
- Moulin-Frier, C., & Oudeyer, P. Y. (2012). Curiosity-driven phonetic learning. In *2012 IEEE international conference on development and learning and epigenetic robotics (ICDL)* (pp. 1–8). IEEE.
- Ngo, H., Luciw, M., Forster, A., & Schmidhuber, J. (2012). Learning skills from play: Artificial curiosity on a katana robot arm. In *The 2012 international joint conference on neural networks (IJCNN)* (pp. 1–8). IEEE.
- Nguyen, M., & Oudeyer, P.-Y. (2013). Active choice of teachers, learning strategies and goals for a socially guided intrinsic motivation learner. *Paladyn Journal of Behavioural Robotics*, 3(3), 136–146.
- Oudeyer, P. Y. (2014). Socially guided intrinsic motivation for robot learning of motor skills. *Autonomous Robots*, 36(3), 273–294.
- Oudeyer, P.-Y., & Kaplan, F. (2006). Discovering communication. *Connection Science*, 18(2), 189–206.
- Oudeyer, P.-Y., Kaplan, F., & Hafner, V. V. (2007). Intrinsic motivation systems for autonomous mental development. *IEEE Transactions on Evolutionary Computation*, 11(2), 265–286. doi:[10.1109/TEVC.2006.890271](https://doi.org/10.1109/TEVC.2006.890271).
- Oudeyer, P. Y., Baranes, A., & Kaplan, F. (2013). Intrinsically motivated learning of real-world sensorimotor skills with developmental constraints. In G. Baldassarre & M. Mirolli (Eds.), *Intrinsically motivated learning in natural and artificial systems* (pp. 303–365). Berlin Heidelberg: Springer.
- Pape, L., Controzzi, C. M. O. M., Cipriani, C., Foerster, A., Carrozza, M. C., & Schmidhuber, J. (2012). Learning tactile skills through curious exploration. *Frontiers in Neurobotics*, 6(6), 11–16.
- Patočka, J. (1998). *Body, community, language, world* (E. Kohák, Trans.). Chicago and La Salle: Open Court.
- Ratcliffe, M. (2008). *Feelings of being: Phenomenology, psychiatry and the sense of reality*. Oxford: Oxford University Press.
- Rietveld, E. (2008). The skillful body as a concerned system of possible actions: Phenomena and neurodynamics. *Theory & Psychology*, 18(3), 341–363.
- Rietveld, E., & Brouwers, A. A. (2016). Optimal grip on affordances in architectural design practices: An ethnography. *Phenomenology and the Cognitive Sciences*,. doi:[10.1007/s11097-016-9475-x](https://doi.org/10.1007/s11097-016-9475-x).
- Rietveld, E., Denys, D., & van Westen, M. (2017). Ecological-enactive cognition as engaging with a field of relevant affordances: The skilled intentionality framework (SIF). In A. Newen, L. de Bruin, & S. Gallagher (Eds.), *Oxford handbook of cognition: Embodied, enactive, embedded and extended*. Oxford: Oxford University Press.
- Rietveld, E., & Kiverstein, J. (2014). A rich landscape of affordances. *Ecological Psychology*, 26(4), 325–352. doi:[10.1080/10407413.2014.958035](https://doi.org/10.1080/10407413.2014.958035).
- Schmidhuber, J. (2010). Formal theory of creativity, fun, and intrinsic motivation (1990–2010). *IEEE Transactions on Autonomous Mental Development*, 2, 230–247.
- Schoupe, N., Braem, S., Houwer, J. D. A., Verguts, T., & Ridderinkhof, K. R. N. (2014). No pain, no gain: Affective valence of congruency conditions changes following a successful response. *Cognitive, Affective & Behavioural Neuroscience*, 15(1), 251–261.
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, 275, 1593–1598.
- Schwartenbeck, P., Fitzgerald, T., Dolan, R. J., & Friston, K. (2013). Exploration, novelty, surprise, and free energy minimization. *Frontiers in Psychology*, 4, 710. doi:[10.3389/fpsyg.2013.00710](https://doi.org/10.3389/fpsyg.2013.00710).

- Schwartenbeck, P., FitzGerald, T. H. B., Mathys, C., Dolan, R., & Friston, K. (2014). The dopaminergic midbrain encodes the expected certainty about desired outcomes. *Cerebral Cortex*, doi:[10.1093/cercor/bhu159](https://doi.org/10.1093/cercor/bhu159).
- Seth, A. K. (2015). The cybernetic Bayesian brain from interoceptive inference to sensorimotor contingencies. In T. Metzinger & J. M. Windt (Eds.), *OpenMIND: 35(T)*. Frankfurtam Main: MIND Group. doi:[10.15502/9783958570108](https://doi.org/10.15502/9783958570108).
- Smith, L. B., & Thelen, E. (2003). Development as a dynamic system. *Trends in Cognitive Sciences*, 7(8), 343–348.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- Thompson, E. (2007). *Mind in life: Biology, phenomenology, and the sciences of mind*. Cambridge, MA: Harvard University Press.
- Tetlock, P. E., & Gardner, D. (2016). *Superforecasting: The art and science of prediction*. New York: Broadway Books.
- Van de Cruys, S. (2017). Affective value in the predictive mind. In T. Metzinger & W. Wiese (Eds.), *Philosophy and predictive processing* (Vol. 24). Frankfurt am Main: MIND Group. doi:[10.15502/9783958573253](https://doi.org/10.15502/9783958573253).
- Van de Cruys, S., & Wagemans, J. (2011). Putting reward in art: A tentative prediction error account of visual art. *I-Perception*, 2(9), 1035–1062. doi:[10.1068/i0466aap](https://doi.org/10.1068/i0466aap).