Ian Robertson[1] and
Michael D. Kirchhoff[1]

Anticipatory Action

Active Inference in Embodied Cognitive Activity

Abstract: *This paper addresses the cognitive basis of anticipatory action. It does so by taking up what we call the acuity problem: the problem of explaining how skilled action seems, on the one hand, to be executed and unfold automatically and reflexively and, on the other hand, to involve anticipation of context-sensitive and constantly changing conditions in performance. The acuity problem invites two contemporary forms of reply, which we label non-inferential enactivism and Helmholtzian inference, respectively. We advance a third avenue for replying to the acuity problem, which takes active inference under the free energy principle as its theoretical starting point. This third way is, we contend, preferable to the other two across a number of important theoretical dimensions.*

Keywords: acuity problem; anticipation; Helmholtz; predictive processing; active inference; enactivism; free energy principle; model; non-inferentialism.

Correspondence:
Email: ianrob@uow.edu.au

1 Department of Philosophy, Faculty of Law, Humanities and the Arts, University of Wollongong, Wollongong, Australia.

1. Introduction

There is a philosophical puzzle at the heart of our understanding of skilled action. To see this, consider a baseball outfielder running to intercept a fly ball. On the one hand, catching a fly ball is a skilled action that seems almost automatic. Well-honed skilful action seems to be executed and unfold automatically and reflexively, ad certainly without the involvement of deliberation or reflection.[2] Indeed, deliberation on the part of the outfielder would seemingly be too slow, attention-hogging, and effortful to properly account for this kind of on-the-fly, super-fast activity. In a slogan: skilful performers do not consider their actions, they just do them. On the other hand, skilled activity cannot simply be automatic and unreflective. Skilled actors are often characterized by virtue of their demonstrating such astounding control over their motor acuity, and by their having honed a capacity to anticipate and adapt to unfolding events on the fly and in ways indicative of their being intricately attuned to present context. Outfielders in baseball, for example, intelligently adjust strategy, often rapidly, and often under highly-pressuring circumstances, in the heat of the moment. We will refer to this puzzle as the *acuity problem* henceforth.

There are several approaches to understanding the cognitive basis of anticipatory action. A classical approach casts the skilled performer as transducing information from different sensory receptors into rich and reconstructive models of the real-world scene, and then using those models to guide their actions. Said approach rests on what Brooks (1999) called the 'sense-model-plan-act' view of cognitive architecture. The problems with this approach are familiar. It treats the brain as passively 'waiting' to be activated by perception. Yet given the presence of sensory and motor processing delays, anticipating sensory input ahead of time turns out to be central in enabling swift and acute performance (Hayhoe *et al.*, 2012). The sense-model-plan-act framework, functioning to reconstruct the world using stored knowledge, has therefore been deemed unfit to account for anticipatory activity prior to impressions by the world on the sensorium. Indeed, such views have been repeatedly exposed as failing to acknowledge the dynamically-updated, on-the-fly character of the

2 David Papineau captures this sentiment nicely in his discussion of cricket batting, when he claims that there is 'no time to think when the ball is released. You can only react' (2013, p. 177; see Fridland, 2017, and Montero, 2016, for criticism of this view).

cognition involved in skilful action (Hurley, 2001; Hutto and Myin, 2017; Sutton *et al.*, 2011).

It is becoming increasingly clear that to make headway on the acuity problem the once-standard picture needs to be flipped on its head. Instead of casting the brain as idly ‘waiting’ on sensory input in order to compute a plan of action for catching a fly ball, say, the brain is now increasingly viewed as serving up anticipations prior to sensory input that approximate the unfolding of worldly and bodily states (Bar, 2007). The cognitive activity involved in skilful action, then, is here depicted as anticipatory in such a way as should not be characterized as entirely an automatic, reflexive, and reactive response. And yet we should also resist its being cast as involving any labour-intensive, time-consuming processes of decision making conditioned on abductive reasoning.

In this paper we consider two contemporary and competing ways of understanding the cognitive basis of anticipatory action.

The first we call *non-inferential enactivism* (Gallagher, 2017; Hutto and Myin, 2017). This brand of enactivism aims to explain anticipation without any appeal to inferential processes.[3] It not only resists the idea of skilled action involving slow and time-consuming inferences (such as those associated with inferring a proposition Q from premises P and $P \rightarrow Q$), but goes further in stating that inferential processes are not part of the explanatory basis of anticipation. On this account, anticipation is treated as a property of the entire cognitive system, which is taken to be a distributed system comprising elements of brain, body, and world. For enactivists of this stripe, anticipating the speed and trajectory of the fly ball results from internal (neural, bodily) dynamics becoming increasingly attuned to the external dynamics of the local environment, where these attunements are taken to be non-inferential in character.

The second we dub *Helmholtzian anticipation* (Hohwy, 2013; Kiefer, 2017), as it conceives of anticipation in terms of Helmholtz’s notion of unconscious perceptual inference. This view is one way of thinking about inference in the now very influential predictive processing scheme in theoretical neuroscience and philosophy. According to Helmholtzian anticipation, running and catching a fly ball is

[3] We make use of the term ‘non-inferential enactivism’, as opposed to merely enactivism or radical enactivism, to capture any enactive view that eschews talk of inferential processes at the explanatory basis of cognition.

possible because brains like ours make probabilistic inferences about the causes of sensory input, enabling a player to latch onto salient environmental properties, as indicated by a forward flow of prediction error,[4] whilst continuously tuning his or her model of the world in light of new incoming evidence from the sensory periphery.[5]

We propose to go between the horns of these two approaches. We will argue that anticipation is a feature of sensorimotor dynamics coupling agent to environment. We find a reason in support of this conclusion in the active inference framework, from Section 3 and onwards. We agree with non-inferential enactivism that anticipation is a property of agent–environment couplings. Yet we disagree that anticipation is non-inferential. We shall develop a view of anticipation as inferential by addressing the acuity problem through the lens of what has been termed *Bayesian enactivism* (*cf.* Allen, 2018), which we associate with the active inference framework under the free energy optimization principle (Friston, 2018; Kirchhoff *et al.*, 2018; Ramstead, Kirchhoff and Friston, 2019). This allows us to construct a compromise between non-inferential enactivism and Helmholtzian anticipation. Active inference differs from Helmholtzian anticipation in the following four ways:

1. Helmholtzian anticipation falls under the predictive coding scheme (Rao and Ballard, 1999) and the Bayesian brain hypothesis (Knill and Pouget, 2004). Active inference is, however, not simply a view of the nervous system as reducing prediction error through an action-based form of perceptual inference. It treats prediction error minimization as realized in the active sampling of the environment by an embodied agent (Friston, 2018).
2. Helmholtzian anticipation implies that anticipation is a property of the internal states of an agent such that the agent encodes a model used to anticipate future events in the world. Active inference casts both perception and action as involved in prediction error minimization, making the entire inferential

[4] Prediction error is a quantity that refers to a mismatch between a *prior expectation* about the causes of sensory input and the actual input.

[5] Model here refers to a generative model, which simplistically speaking is a model that generates predictions about the causes of sensory input and which can be updated by taking into account new evidence (i.e. by working to cancel out any backward flowing error signals).

processes between action and perception integrated and embodied, with inference spanning brain and body, on the one hand, and action dynamics, on the other. Thus, inference on this model is not 'purely internal' (Pezzulo, Kemere and Van Der Meer, 2017; see also Allen *et al.*, 2019; Friston *et al.*, 2012).

3. Helmholtzian anticipation takes anticipation to be representational, assuming that it is not possible to forecast events in the absence of states representing such events. We will argue that under active inference there is no need necessarily to posit the presence of representations to provide an account of anticipation to address the acuity problem.
4. Helmholtzian anticipation typically casts internal (brain) states as informationally secluded from external (worldly) states. Active inference, however, implies an informational coupling between internal and external states given sensory and active states (Kirchhoff and Kiverstein, 2019b; Ramstead, Kirchhoff and Friston, 2019). Given restrictions of length, we omit further discussion of this specific point.

The structure of this paper is as follows. In Section 2 we review key tenets of non-inferential enactivism and Helmholtzian anticipation, and explicate how they differ in explaining the cognitive basis of anticipatory action. In Section 3 we suggest that a plausible view of anticipatory action can be located in between non-inferential enactivism and Helmholtzian anticipation. We shall argue that the mechanism enabling anticipation is a species of inference, which we will cast in terms of active inference. We will highlight the plausibility of casting anticipation in terms of active inference by demonstrating it to be significantly distinct from Helmholtzian anticipation (points 1–4 above), while, at the same time, addressing points of overlap and contrast with non-inferential enactivism. In Section 4 we turn to consider two objections to our inferential characterization of anticipatory action. We will argue that both objections rest on misconceptions about the properties of active inference. We end this paper by showing that understanding anticipatory action as inferential — *qua* the active inference framework — can help in developing a novel set of first steps towards addressing the acuity problem. Such a characterization, we submit, allows us to explain the intelligent aspects of the anticipation involved in skilful action in inferential terms. That said, it does not require characterizing anticipation as a form or variation of Helmholtzian inference. In this way, we are concerned with the following research question: in what sense might active inference, by way of walking a sensible line between inferentialism and non-inferentialism, provide a framework within which to approach the acuity problem?

2. The Outfielder Problem: Anticipation versus Inference

Our aim in this section is to examine how non-inferential enactivism and Helmholtzian anticipation purport to explain the cognitive basis of anticipatory action. We will do this by highlighting salient differences between them in how they explain the cognition involved in anticipation, before showing how these differences result in their offering competing solutions to the outfielder problem (the problem of explaining how expert baseball outfielders are so gifted at anticipating and thereby intercepting fly balls).

2.1. Helmholtzian anticipation

One solution to the acuity problem casts the cognitive basis of anticipatory action as a form of unconscious inference, where inference is understood as a species of inference to the best explanation realized by neuronal activity. This view is typically defended under the auspices of the predictive processing theory of mind (PP). PP states that brains must deal with uncertainty about the (hidden) causes of sensory observations, and claims that the best way for the brain to handle uncertainty is by approximating a form of Bayesian inference. This is not an inference in the conscious sense of the word; rather, it is unconscious and probabilistic, arguably driven by conformity with Bayesian norms.[6]

Proponents of PP take one of the main explanatory advantages of their theory to be that it explains how the brain solves an inverse problem (Hohwy and Michael, 2017). The key assumption is that the brain is taken to be epistemically isolated from the world, and thereby 'needs to represent the world so we can act meaningfully on it, that is, it has to figure out what in the world causes the sensory signal it

[6] There is ongoing and increasingly intricate debate as to the sense in which a cognitive system would need to be constrained by or comply with Bayesian norms in order to be characterized as non-trivially performing a Bayesian operation (see e.g. Orlandi, 2016; Colombo, Elkin and Hartmann, 2018). We mention this here in order to set it aside — a task for a different paper.

receives' (Hohwy, 2012, p. 2; see also Clark, 2013). The inverse problem demands explanation as to how the brain infers the (hidden) causes (or latent variables) of the sensory input (or data) it receives. The difficulty is that impinging input is assumed to be ambiguous, uncertain, and noisy. Hence, the sensory stream underdetermines its hidden (true) causes. PP provides an answer to the inverse problem by claiming that the brain uses hierarchical generative models in approximating a probabilistic form of causal inference that minimizes prediction errors between actual and expected sensory signals. In this fashion, generative models function so as to probabilistically map hidden causes to sensory data under prior expectations.

The idea of unconscious inference underwriting perceiving and cognizing finds a precursor in the work of Hermann von Helmholtz (1867/1962). He compared perceptual inference to syllogisms, the premises of which have to be garnered and established by means of induction. PP, however, does not characterize the brain as engaged in inference of a traditional form, whereupon there is a transition between premises and a conclusion. Instead, the updating of neurally-instantiated predictive models is taken to be inferential in the sense of approximate Bayesian inference. Thus, what we perceive is the expression of a form of abductive inference under prior (Bayesian) beliefs. This means that inference is a matter of sampling evidence for prior expectations (i.e. the brain's model).

Cognizing (including perception) on this view is just a matter of reducing prediction error, which can be cast as garnering evidence for one's prior beliefs. According to models of PP predicated on Helmholtzian inference, this can be achieved in one of two ways: (a) via perceptual inference — i.e. updating one's model to fit with incoming sensory observations; or (b) via active inference — i.e. to act so as to make the world fit with one's expectations. By adopting the latter strategy, Hohwy thinks that the brain 'vicariously enslaves an external body plant to fulfill its predictions of sensory input, given its internal model of the world and itself' (2016, p. 276). By making use of the former, the brain is conceived of as testing hypotheses in a manner akin to a scientist. As Burr and Jones put it: 'Just as scientists test hypotheses by conducting experiments using well-calibrated lab equipment, so too perceivers must test their predictions by using their bodies to interact with their environment' (2016, p. 596; see Bruineberg, Kiverstein and Rietveld, 2018, for discussion).

2.2. Non-inferential enactivism

Enactivists find inferential views of mind, including Helmholtzian PP, to be problematic. They think that casting the brain as engaged in a form of unconscious inference is explanatorily otiose (at best), if not outright theoretically distorting (at worst). Enactivists emphasize the role of the acting body in generating cognition. The fundamental enactivist credo is that cognition is realized not only by neural goings-on but in a dynamic network comprising brain, body, and world (Varela, Thompson and Rosch, 1991; Thompson, 2007; Colombetti, 2014; Di Paolo, Buhrmann and Barandiaran, 2017; Hutto and Myin, 2013; 2017). Given that the brain is part of a system including parts of the local microenvironment, enactivists take their view to undermine the rationale for claiming that the brain need execute complex inferences to 'get at the world'. Hence, there is no reason to think that brains must do work inferring a model of the world so that the organism can then act upon the model in order to act in the world (Hutto, 2018).

Gallagher (2017) has provided a recent defence for what we here call non-inferential enactivism. For Gallagher, construing the brain as employing inference in order to establish epistemic access to the world ignores the fact that perception is highly organism-relative, and how an organism's perception is determined by its anticipation of reward — its 'ulterior motives' (*ibid.*, p. 116). He argues that positing complex, neurally-realized inference to explain how the brain overcomes such isolation threatens to paint perception as a straightforwardly intellectual process, the role of which is 'identifying or recognizing objects or guiding bodily movement in the world' (*ibid.*; see also Varela, Thompson and Rosch, 1991, p. 136). Non-inferential enactivism recognizes that the organism does not receive 'sensory data first, followed by inferential processes that conclude reward possibility (an additional neural or cognitive function added to sensory activation)' (Gallagher, 2017, p. 116). It instead construes perception as 'an intrinsically reward-oriented response or attunement to stimuli due to prior experiences and plastic changes — there's *no room for or need for inference*s in this respect. Perception is already attuned to reward possibilities' (*ibid.*, emphasis added).

Non-inferential enactivism — following research principles established by ecological psychology — also takes the notion of the brain receiving only impoverished sensory data from the world (and thereby having to 'solve' its inverse problem by engaging in complex inference) to be undermined by the exploratory nature of organisms, and how they are almost always moving around and investigating their environmental niche. Gallagher advances exactly this view when he writes:

> The poverty of stimulus problem… is addressed by the possibility of bodily movement… Moving around the environment provides more information and reduces the ambiguity. The point in such action is that the environment specifies itself (the environment is what it is) — it is not impoverished; the poverty only arises if we think that the brain has no access to the rich structure of the environment. It disappears if we acknowledge that the organism has access — is attuned or coupled — to the environment over time and is not only capable of movement, but is almost always moving. (*ibid.*, p. 119)

Thus, non-inferential enactivists reject the idea that the brain receives for the most part noisy and ambiguous input. Due to the organism's capacity to investigate and probe its environment (changing it around its own requirements along the way) while navigating it, the brain will not, for the most part, receive only highly underdetermined sensory signals. As Froese and Ikegami (2013) put it, the brain should be characterized 'not as a black-box prison of the mind, but rather as a self-organized perspectival reference point that serves to enact a set of meaningful relations with its milieu' (p. 33; see also Dreyfus, 2002; Orlandi, 2014). *Qua* non-inferential enactivism, then, recognizing the propensity of the organism to be constantly investigating and interacting with its environment undercuts the motivation for arguing that brains perform unconscious inference.

2.3. The outfielder problem: Helmholtzian and non-inferential solutions

Helmholtzian anticipation and non-inferential enactivism offer competing solutions to the outfielder problem. The outfielder problem can be summarized, briefly, as follows. Professional baseball outfielders demonstrate exceptional skill when pacing their run so as to position themselves suitably for intercepting fly balls. Empirical data, however, indicate that stationary outfielders are not very adept at predicting the landing locations of fly balls (Shaffer and McBeath, 2005). Explanation is required as to how outfielders are — while running to make a catch — able to anticipate the trajectory the fly ball will take and dictate their locomotion accordingly. A solution to the outfielder problem, then, would need to answer the following question: how are running outfielders able to anticipate where and when to end up so as to successfully intercept fly balls?

The non-inferential enactivist solution to the fly ball problem can be viewed as a response to more traditional accounts, on which it is conjectured that expert outfielders, in their pursuance of intercepting a fly ball, construct and utilize an internal predictive model for the purposes of inferring its landing location, and guide their run accordingly. This kind of traditional solution to the outfielder problem depends on the outfielder's constructing and utilizing an internal model of the fly ball trajectory, maintaining and updating it even while running to make a catch. Said model would have to be sensitive to an extremely intricate and interwoven manifold of factors. Factors such as how wind direction is likely to affect the ball's flight path, and other complex (often random) variables.[7]

A non-inferential enactivist solution to the outfielder problem differs from this kind of traditional, inferential one primarily in that it jettisons any appeal to the notion of the outfielder exploiting a detailed internal reconstruction or representation of the fly ball trajectory (Anderson, 2014; Gallagher, 2017). Instead, non-inferential enactivists take the outfielder to be executing what has come to be known as the Chapman strategy. The Chapman strategy, named after noted physicist

7 One way to cast this distinction between inferential and non-inferential accounts of the outfielder problem, specifically, or action, more generally, is in terms of a distinction between model-based and model-free strategies to choice and decision making. Crudely put, model-based strategies are ones where action results from deployment of rich and context-sensitive encoded knowledge. Model-free strategies, on the other hand, turn on learning about the world via trial and error, without invoking a model of the world (see e.g. Stepp and Turvey, 2009). A simplified picture of the distinction between these strategies would seem to map onto familiar, yet almost certainly too simplistic, divisions between habit and reason, or automaticity and intelligence (*cf.* Fridland, 2017). This overly simplistic division between model-based and model-free strategies assumes rather than addresses the features giving rise to the acuity problem. Looking ahead, the active inference framework provides a unified theoretical basis for 'delicately combining [these] two modes [i.e. strategies] within an overarching economy, adaptive agents may identify the appropriate contexts in which to deploy the model-free ("habitual") schemes. "Model-based" and "model-free" modes of valuation and response, if this is correct, simply name extremes along a single continuum and may appear in many mixtures and combinations determined by the task at hand' (Clark, 2016, p. 255). Active inference also implies that there are no clear-cut boundaries between these two strategies, as most agents are capable of both, often in a combined form (*ibid.*). There is an ongoing discussion in the literature as to whether the model-based vs. model-free distinction maps onto the representation vs. non-representation distinction in forms of cognition under active inference. Addressing this issue will be a task for another occasion.

Seville Chapman, says that expert outfielders run under the fly ball — increasing and decreasing acceleration and changing running direction where appropriate — in such a way as to nullify or abrogate the perceived vertical acceleration of the fly ball (Chapman, 1968; see also Fink, Foo and Warren, 2009; Dienes and McLeod, 1993; Postma *et al.*, 2017; see also Beer, 2000; Anderson, 2014, pp. 184–5). Her doing so will, so contended Chapman, reliably result in the skilful outfielder being well-positioned to intercept the fly ball.

Characterizing the outfielder as performing the Chapman strategy, then, dispenses with any explanatory need to appeal to the notion of the outfielder employing an internal predictive model for representing the arc and trajectory or inferring the landing location of the fly ball. As Fink, Foo and Warren (2009) put it, 'fielders are led to the right place at the right time by coupling their movements to visual information in a continuous "online" manner' (p. 1).[8]

In this way, a Chapman solution to the outfielder problem apparently does away with the need to characterize the outfielder as engaged in any rigorous, labour-intensive inference, or detailed, cognitively-exorbitant computations of the fly ball's trajectory (Shapiro and Spaulding, 2018; Anderson, 2014; Clark, 1997). As Gallagher puts it, there is, on this non-traditional account of outfielder catching, no necessity for the outfielder 'to compute in-the-head mental representations — of the ball, its speed, its trajectory, and so on' (2017, p. 14). The non-inferential enactivist takes this to be an example of anticipatory action that does not require inference. The outfielder 'is operating within the situation itself rather than on a model of the situation inferred by the brain' (*ibid.*, p. 115).

Although more traditional, inferential solutions to the outfielder problem are increasingly unpopular among contemporary sports psychologists, one might intuitively take such solutions to be consistent if not heavily consonant with Helmholtzian anticipation. After all, Helmholtzian anticipation — in the contemporary PP guise we have been considering — posits hierarchical probabilistic models that unfold across multiple spatial and temporal scales. The advantage of the brain employing these models is that they enable us, as Clark puts it, 'to lock us onto worldly causes that are ever more recondite, capturing regularities visible only in patterns spread far in space and

---

8 What we refer to here as the Chapman strategy is often referred to in the current sports psychology literature as the optical acceleration cancellation strategy.

time' (2015, p. 5). One might thereby take PP to suggest that the outfielder is endowed with the capacity to model or represent exactly the kinds of exceptionally complex and concealed variables that would need to be calculated in order to infer the likely landing location of the fly ball. Granted, such a solution might seem to introduce an unsettling divide between the outfielder and her environment, but, as we have already seen, some proponents of Helmholtzian anticipation seem all too happy to allow for 'a schism between the prediction-generating models of the brain and the modeled states of affairs in the world' (Hohwy, 2016, p. 5).

As is happens, though, the most prominent advocates of Helmholtzian anticipation — even those who allow for a clear divide between mind and world — anticipate the landing location of the fly ball *by virtue of* constructing and operationalizing an explicit, detailed predictive model of its flight trajectory. In fact, they take PP to support the idea that skilled outfielders will carry out the Chapman strategy: they argue that the outfielder could 'simply treat as salient (highly weighted) all and only the prediction errors associated with optical acceleration of the ball', and then run in order to nullify perceived optical acceleration (Clark, 2017). The outfielder's strategy of running to do away with perceived optical acceleration will then amount to their suppressing prediction error (prediction errors will be engendered whenever there is perceived optical acceleration). In this way, PP advocates deny that the outfielder will need to 'rely on internal inference', such that she is required 'to compute the arc, acceleration, and distance of the ball much like one would do by following a physics textbook' (both quotes from Hohwy, 2016, p. 278). Yet what is not up for discussion is that inference in the form of approximate Bayesian inference is centrally involved. Even if the use and sampling of sensory input may seem to be less substantial in cases such as fly ball catching, 'the kind of selective, sparse sampling in play here requires heavy, explicit modeling of external causes, including complexity reduction and updating of expected precisions' (Hohwy, 2016, p. 279). For this reason, Hohwy claims that skilled action requires the brain to harness 'a model that is rich, slow and reconstructive in the sense that it has significant causal depth, which it uses to take causal interactions at multiple spatiotemporal scales into account when generating predicted sensory flows' (Hohwy, 2019, p. 5). We will seek to disabuse any persuaded readers of this notion in Section 3.3.

3. Anticipation as Active Inference

Our aim in this section is to show that it is possible to go between the pillars of Helmholtzian anticipation and non-inferential enactivism. We argue the notion of unconscious inference posited by Helmholtzian anticipation is problematic, but, at the same time, claim that there is a significant sense in which anticipatory action is inferential. As mentioned in Section 1, we will do this by unpacking inference via the *active inference* framework (AIF) — also known as the free energy optimization principle (Friston *et al.*, 2017). Approaching anticipation from this perspective offers a set of solutions to the acuity problem that neither Gallagher's form of non-inferential enactivism nor Helmholtzian anticipation can properly address.

3.1. The brain–body–environment system

Gallagher (2017) contends that we should prefer the non-inferential enactivist solution to the outfielder problem to the Helmholtzian alternative. For him, the assumptions underlying Helmholtzian anticipation are clear yet problematic. As we saw above, one assumption is that the brain is hidden away inside the skull, having only access to its own predictions and error signals. Crucially, such a view implies 'the brain has no direct access to the world' (*ibid.*, p. 109). Another assumption is that all cognition (including perception) is realized in the brain. In the parlance of philosophy of mind this implies realizer internalism, i.e. cognitive activity is taken to be realized exclusively by neural activity. The other explanation is enactivism. The latter explanation turns on different assumptions and is superior, Gallagher claims, because:

> The human brain not only evolved along with the human body, and works the way it does because of that; it's also not isolated, but rather is dynamically coupled to the body that is dynamically coupled to an environment. The organism (the brain–body system) is operating within the situation itself rather than on a model of the situation inferred by the brain. This coupling of brain–body–environment is structured by the physical aspects of neuronal processes, bodily movements, affects, anatomy and function, and environmental regularities. (*ibid.*, p. 115)

Non-inferential enactivists find this latter explanation, or something close to it, a convincing reason to be sceptical about optimization schemes such as predictive processing, and the AIF (*ibid.*; Hutto and Myin, 2017). Gallagher sketches a view in which cognitive activity is not simply a result of brain functioning; rather, it arises out of a complex and delicate interplay between resources spread across brain, body, and world (for further detail, see Sections 2.2 and 2.3 above). On the enactive account, there is no need to appeal to Bayesian inference and principles from machine learning to explain cognitive activity (we return to this point in the next subsection).

We find Gallagher's dismissal of the AIF too quick. The reason is that the AIF can accommodate the enactivist view of cognitive systems being extended in the sense of extended organism–environment systems coupled via perception and action (Kirchhoff, 2015; 2017; 2018; Ramstead, Kirchhoff and Friston, 2019).

Under active inference, organisms act to minimize variational free energy (Friston, 2010). The *active* element in active inference captures the idea that agents work to infer action policies to ensure that their actions align with their expectations about sensory input. In other words, in active inference agents not only act to reduce free energy (the sum of prediction error), they also act to minimize *expected free energy*, i.e. the free energy associated with the sensory state they will occupy conditioned on future action. In this sense, active inference captures the counterfactual character involved in recurrent cycles of perception and action; or, in sensorimotor contingencies (Mirza *et al.*, 2016; Seth, 2014). Technically, variational free energy is a mathematical bound on 'surprisal' or 'surprise' (Friston *et al.*, 2012; Kirchhoff *et al.*, 2018). The time average of surprise is 'entropy', which in information theory is a measure of uncertainty. In the AIF, agents reduce surprise by minimizing variational free energy.

The connection between variational free energy and surprise is important. Surprise amounts to a measure of how surprising it would be for a system to inhabit a specific state given the kind of organism it is. This suggests that surprise is conditioned on the phenotypic states of an organism. Crucially, as Hohwy observes, surprise 'cannot be changed by perceptual inference because perceptual inference changes the hypotheses about the sensory input and not the sensory input itself' (2013, p. 85). Active inference is a story about how it is possible to change the sensorium by changing the hidden or external causes of sensory input. It thus takes centre stage in accounting for how agents minimize surprise. Under active inference, organisms infer the policy (i.e. action routine) that is most likely to combat an increase in surprise. For example, if one is hungry, one policy to select for would be to open the kitchen fridge and get to cooking. Yet it will not suffice merely to be inferring over a distribution of action policies, if the aim is to satisfy one's hunger, say. Embodied action is needed to align prior beliefs about expected sensory observations with actual sensory observations. Crudely put, locking your front door is more likely to result in you being in states with low surprise, avoiding the highly surprising state of being confronted by an intruder (Chen *et al.*, 2019). This active and embodied dimension of active inference is precisely that sensory outcomes at any given time depend on hidden states or causes, while such states, in turn, evolve in a way that depends on action (Mirza *et al.*, 2016). The AIF is thus a story about reducing expected free energy via inferring action policies on the side of the generative model, on the one hand, and ensuring that expected sensory inputs are met via embodied action in the local environment, on the other hand.

It is unsurprising that Gallagher takes something like predictive processing *qua* Helmholtzian inference as implying realizer internalism, for cognition on this view is usually cast as the result of *Bayesian belief updating* performed over neuronal states, i.e. inferring *posterior probabilities* of predictions, P(H|E), conditioned of a likelihood function, P(E|H), and a prior probability, P(H). Helmholtzian anticipation is in this sense akin to other theories such as predictive coding and the Bayesian brain hypothesis. Yet this group of views has been argued to be incomplete theories of how we infer states of the environment. As Friston has put it: the 'missing bit is the *enactive* compass of the [active inference framework]. In other words, the [AIF] is not just about making the best (Bayesian) sense of sensory impressions of what's "out there". It tries to understand how we sample the world and author our own sensations' (Friston, 2018, p. 12). As we saw above, perceptual inference *per se* cannot address how an agent is able to minimize surprise. If this is correct, it follows that it is to active inference one must turn to find an account of surprise minimization. This also means that the AIF denies that it is possible to engage in model optimization exclusively via perceptual inference, i.e. by working only to optimize one's posterior (Bayesian) beliefs via perception.

In the AIF, authoring our own thinking and sensations is conditioned on *enacting a generative model* (Ramstead, Badcock and Friston, 2018; Ramstead, Kirchhoff and Friston, 2019; Kirchhoff and Kiverstein, 2019a). Technically, a generative model is a statistical mapping from hidden (external) causes to input under a set of joint probability distributions of outcomes or consequences and their causes (Parr, Rees and Friston, 2018). If one approaches generative models from a long-term evolutionary perspective, the AIF treats the *entire* *embodied organism* as a model of the organism-relevant dynamics of its niche (Linson *et al.*, 2018, p. 2). Models are in this sense not *distinct* from the organism. Although both AIF and non-inferential enactivism place heavy explanatory emphasis on organismic action, Gallagher worries that schemes like the AIF give some special privilege to the brain in cognition and action. However, as Linson *et al.* (2018) further note: during later evolutionary stages when neural systems arises, 'brains come to augment the more fundamental embodied agent with a neuronal-connectivity-based extension to the generative model that handles more complex organism/niche dynamics… the human (neural) instantiation of active inference… the brain should be understood as "taking a back seat" to the body, serving the body by facilitating more complex coordination' (*ibid.*, p. 2). Brains, on this view, are nested within bodies, which are nested within environments. All factors play a role in action coordination (Clark, 2016; Kirchhoff and Kiverstein, 2019a,b).

This brings out a further point of contrast between predictive processing under Helmholtzian inference and the AIF. The former takes organisms to *have a generative model* — realized over neuronal dynamics. Crucially, under the AIF, the organism does not have a generative model as if there were two distinct things; namely, the organism and then a model inside its head. Instead an organism is understood as *being a model* of its milieu, where being a model is defined as follows: 'We must here understand "model" in the most inclusive sense, as combining interpretive dispositions, morphology, and neural architecture, and as implying a highly tuned "fit" between the active, embodied organism and the embedded environment' (Friston *et al.*, 2012, p. 6). Kirchhoff and Kiverstein (2019b) extend on this, noting that this embodied and enactive characterization of the generative model means that an organism can be cast as a hierarchically nested probabilistic model conditioned on the sensory, physiological, and morphological states that are highly reliable given the kind of life it leads and the environment it inhabits (*cf.* Friston, 2010; 2013).[9]

'Being a model' in this sense is consistent with Conant and Ashby's (1970) *Good Regulator Theorem*, which states that two systems are

---

9 We note that Gallagher (2017, p. 130) refers to the distinction between 'having a model' and 'being a model'. Yet he does not elaborate on it. This seems to us to be a missed opportunity.

coupled to one another when one of the systems can remain in states of low surprise despite pressures on the system from the outside (i.e. its environment). This yields a clear connection to the free energy principle, which states that 'adaptive systems minimize a limit on free energy (long-run average surprise) by inducing and refining a generative model of the causes of sensory signals' (Seth, 2014, p. 8). This means that generative models can be cast as control systems in the precise sense of being 'good regulators' of the larger system (i.e. environment) within which the organism is embedded and with respect to which it must resist perturbations to maintain the surprisal of its states to a local minima.

This speaks to skill and habit formation, elements crucial to successfully catching a fly ball. Recent work by Chen *et al.* (2019) addresses this by introducing a new aspect to the AIF; the ability to update and refine a policy space. First, prior probabilities of sensory input given action are specified over a distribution of possible action policies, where each specific action policy specifies a series of actions over time. This means that learning action policies takes the form of working to optimize or improve the distribution of policies. Second, highly specialized skills can be associated with the acquisition of habits, given policy pruning, i.e. learning that some rather than other policies are ultimately more preferable given the situation. Leveraging the idea of 'being a model' in the sense of the Good Regulator Theorem means that agents, given a history of learning, become good regulators of the kind of environments they tend to frequent. The optimal policy in the case of the running to catch a fly ball is to infer a sequence of actions enabling one to run under the ball — as predicted by the Chapman strategy. Agents prune away, over time, less than optimal policies, thus limiting the space of action possibilities.[10] This kind of policy pruning can account for specialization (behaviours tightly adapted to a given environment or situation), generalization (behaviours requiring the agent to take a larger number of action policies into consideration), and what we might refer to as a kind of mixed-strategy between specialization and generalization. A mixed-strategy is what one might expect is involved in enabling agents to successfully navigate the outfielder problem, for agents need to behave flexibly both to the speed of the ball, direction of wind, and other context-specific contingencies. Policy pruning, as Chen *et al.* go

10 This strikes us as a nice way to address the frame problem in cognitive science as well.

on the observe, 'would explain the effects of practice — as we gain expertise in a given task, the time it takes to complete the task and the subjective experience of planning during the task diminishes, likely because we have learned enough about the structure of the task to discern and learn appropriate habits' (2019, pp. 4–5). Optimizing the distribution over action policies thus speaks to how agents, over time, attune appropriately to the dynamics of their environments via the perception–action cycle, coupling the agent to the environment, and vice versa.

3.2. Active inference isn't just a 'doing'

The claim that adaptive agents are generative models is consist with the enactive emphasis on extended, brain–body–world systems precisely because being a generative model in the inclusive sense as defined above is to be a system dynamically coupled with its local niche. This also falls out of the idea that it is the continuous cycles of perception and action that establishes the conditional independencies between the internal states of the agent and the external states of the environment (Ramstead, Kirchhoff and Friston, 2019). So, there is no obvious reason why one should take Gallagher's dismissal of active inference seriously on this point.

Gallagher (2017), however, still thinks there is a deep tension between the AIF and his brand of non-inferential enactivism. The main issue now turns on the reference to inference in the AIF.

Gallagher frames the issue as follows. He starts by saying:

> With respect to PC [predictive coding] models, enactivist views emphasize a more holistic system of brain–body–environment would clearly favor a move away from internalist and intellectualist vocabularies (and conceptions) of 'hypotheses', 'inference', and 'representation' in favor of more embodied terms like 'adjustment', 'attunement', and 'affordance'. Such terms not only simply substitutes for the PC terms; they change the way that we think of the brain's engagement. (Gallagher, 2017, p. 21)

He considers the AIF. Recall that we argued above that active inference says that perception evolved and is for facilitating action. It denies the Helmholtzian notion that action is for vindicating or disproving perceptual hypotheses like an intervening scientist. This is one of the main differences between the AIF and predictive coding schemes such as predictive processing under Helmholtzian inference. Against this view of the AIF, Gallagher says the following: 'active inference is *not* "inference" at all, it's a *doing*, an enactive adjustment, a worldly engagement' (*ibid.*, p. 19, italics added).

We will now argue that the enactivist references to 'doings', 'adjustment', and 'worldly engagement' come up short when having to address the acuity problem. We find it odd that Gallagher would object to anticipation being inferential, as it is difficult to see how reference to things like doings and worldly engagement can account for conditional future states — i.e. the bringing about of possible future states conditioned on sensorimotor contingencies. Note that any plausible answer to the acuity problem will need to explain how skilful performers 'stay ahead of the curve' and thus anticipate, even in the heat of the moment, how slight variable changes would compel or demand intricate adjustments in action. It is on this precise issue that the surprise minimization framework of the AIF becomes relevant for making headway on the acuity problem, with reference to counterfactual inference (or model inversion) and temporally deep generative models.

In interceptive expertise a common observation is that elite players do not keep or align their central foveal gaze with the ball during its flight path. This observation has been made in a range of different sports, from baseball (Hubbard and Seng, 1954; Bahill and LaRitz, 1984), cricket (Croft, Button and Dicks, 2010; Land and McLeod, 2000), table tennis (Ripoll and Fleurance, 1988), to squash (Hayhoe *et al.*, 2012). A recent study by Mann, Spratford and Abernethy (2013) comparing elite and club-level cricket batters found that elite batters used two distinctive eye movement strategies, suggesting that saccadic eye movements play a significant role in interceptive expertise. They found that experts rely on two *predictive* saccades to anticipate (a) the location of the spot where the ball would bounce after ball-release, and (b) the location of the bat–ball point of contact, enabling the batters to direct their gaze at the ball as they hit it. As they state: 'elite batters directed their gaze ahead of the flight-path of the ball immediately prior to bat–ball contact, whereas the gaze of the club-level batters tended to be behind the ball. The elite batters appeared to use a strategy that ensured they could "park" their gaze ahead of the ball so that gaze could "lie-in-wait" for the ball to arrive' (2013, p. 6). Crucially, Mann *et al.* found that predictive saccades employed to anticipate where the ball would bounce occurred earlier as the skill level in the batters increased, 'reflecting a superior ability to predict the future landing point of the ball' (*ibid.*, p. 1).

Assuming these results are correct, how might one explain this without appeal to inference but merely by reference to *doings*? There seems no denying that batters are doing something; they are engaged in an embodied activity in the world. Is that all they are doing, however? Are there reasons to deny that these doings either are or involve a form of inference? How would *doings* — understood in non-inferential terms — address the issue of *conditional outcomes of future states* given action, when batters make use of epistemic actions such as eye movements to anticipate where the ball will bounce prior to it happening?

One option might be to say that 'neural plasticity mitigates to some degree the need to think that subpersonal processes are inferential. The neural networks of perception are set up by previous experience — "set up to be set off"...' (Gallagher, 2017, p. 115). But this cannot be quite right, for elite batters are not simply 'set up' so as to be 'set off' but are able to handle variations in ball speed, direction, and bounce on the fly. Another option might be to say that batters are responding to what the situation *affords*. The concept of affordance is a term of art in the tradition of ecological psychology (Gibson, 1979/1986). Crucially, it is a *relational* notion in the sense that it belongs to the directly coupled organism–environment system. Specifically, it concerns perception of niche-specific information as providing action opportunities relative on an organism's sensorimotor capabilities. There is something clearly problematic with appealing to affordances to do the explanatory work here. One cannot stand in a directly coupled relation to future states, for such states are yet to be present (Linson *et al.*, 2018).[11] Hence, if Gallagher's brand of non-inferential enactivism is to explain the cognitive basis of anticipatory action, then he needs to explain the capacity of skilful performers to adapt to novelty while maintaining the fluidity of their action.

To make headway on this issue, consider that elite batters are not tracking the ball over its flight path; rather, they are using a saccadic

---

[11] Gallagher might argue that the outfielder effectively engaging in the Chapman strategy means that the ball will continue to be perceived as afforded (as long as her running successfully nullifies perceived optical acceleration). The issue here is that outfielders — as observed by Postma *et al.* (2017) — seem often to perceive fly balls that are uncatchable by virtue of their travelling at very fast speeds *as uncatchable* (unafforded) even before they are at full running capacity. Non-inferential enactivists would need to explain this without appeal to the notion that outfielders are engaged in an inferring of conditional future states.

prediction enabling them to anticipate where the ball would bounce prior to ball–bat impact. Prediction refers to estimating future states of a system (Wolpert and Flanagan, 2001). In cases such as interceptive expertise, prediction takes the form of a *probabilistic mapping* of the *counterfactual* relation between perception and action. Interestingly, counterfactual mappings of this kind are at the heart of non-inferential enactivism. Gallagher, for instance, says that we 'perceive things in terms of… sensorimotor contingencies and in terms of what those things pragmatically afford in relation to a body like mine, in the situation' (2017, p. 116). Two things are important to mention here. First, sensorimotor contingencies underline the importance of perception and action in ongoing organism–environment engagement. Second, sensorimotor contingencies *qua* counterfactual probability relations between perception and action underpin the *future-oriented* structure of anticipation: they are probabilistic mappings of future sensory states conditioned on action. Getting at expected sensory observation, like actually grabbing the ball in mid air, is a matter of selecting and executing the right action policy. Prediction is therefore *not just doings*; it is rather a matter of generative models actively inferring policies that if selected would bring about the expected sensory outcomes, and therefore the minimization of any expected variational free energy. Prediction in this context cannot be understood absent the notion of inference, and certainly cannot be understood *qua* AIF as a form of doing that does not involve inference. Such counterfactually equipped generative models are known as temporally deep generative models. As Friston puts it:

> [S]ystems that can grasp the impact of their future actions must necessarily have a *temporal thickness*. They must have internal models of themselves and the world that allow them to make predictions about things that have not and might not actually happen. Such models can be thicker and thinner, deeper or shallower, depending on how far forward they *predict*, as well as how far back they *postdict*, that is, whether they can capture how things might have ended up if they had acted differently. Systems with deeper temporal structures will be better at inferring the counterfactual consequences of their actions. (Friston, 2017, *Aeon*)

The AIF thus provides us with a story about sensorimotor contingencies, cast in terms of probabilistic inference, yet realized by embodied engagement in the world.

Perhaps one could object that Gallagher's (2017) claim is consistent with the claim that active inference involves inference but is not exhausted by it, because there is more to actions or what Gallagher terms doings than inference. The problem with this worry is that there are *no* doings that are not subserved by inference according to the AIF. Gallagher might object, however, that the emphasis on inferential models under the AIF neglects the key role of the action–perception cycle coupling organism and environment dynamics. There is, however, a complication with any such objection; namely, it overlooks the important distinction between generative model and generative process. The generative process couples the organism to its environment, and vice versa, via sensing and acting. Put bluntly, there can be no sensory observations without the generative process. The generative process can therefore be understood in terms of the pragmatic and epistemic actions of an organism coupling organism and environment, where the modulation *via action* of the coupling relation reduces surprise. In this specific sense, the AIF implies the presence of an informational coupling between organism and its local niche (Kirchhoff and Kiverstein, 2019b; Ramstead, Kirchhoff and Friston, 2019).

We ended the previous section by considering the notion of optimizing a distribution of action policies by pruning suboptimal policies via learning, eliciting a view of skill and habit formation under the AIF. We can now add to this account by noting that the action policies being either selected for or pruned away are policies of counterfactual sensorimotor contingencies, over the temporal depth of a hierarchically organized generative model.

3.3. Active inference isn't (necessarily) representational

We now turn to assess and dismiss another reason why Gallagher (2017) thinks there is a deep tension between the AIF and his brand of non-inferential enactivism, showing the latter to be preferable to the former. The issue is going to turn on the claim that inferences are *representational* and therefore incompatible with non-inferential enactivism, which conceives of the vast bulk of cognitive activity in non-representational terms.

Gallagher wants to convince us that rather than invoking the vocabulary of representations, even in their action-oriented specifications (*cf.* Clark and Grush, 1999), we should rather think of cognitive activity as 'a kind of ongoing dynamical adjustment in which the brain, *as part of and along with the larger organism*, settles into the right kind of *attunement* with the environment — an environment that is physical but also social and cultural' (Gallagher, 2017, p. 18, second italics added).

The notion of 'attunement' is important in Gallagher's account. But what does it amount to? It is somewhat surprising to observe that Gallagher only refers to this notion twelve times in his 2017 book. He says for instance that 'social cognition is an attunement process that allows me to perceive the other as someone to whom I can respond…' (*ibid.*, p. 12). He also says that 'brains play an important part in the ongoing dynamical attunement of organism to environment' (*ibid.*, p. 20). Details are lacking. The closest Gallagher comes to spelling out what attunement means is when he states: 'what do brains do as part of a dynamical attunement of organism to environment…? The answer is that the brain participates in a system, along with eyes and face and hands and voice, and so on' (*ibid.*, p. 163). This might be correct, even if somewhat abstract. In what follows we give an active inference account of attunement, giving this notion a firmer information-theoretic articulation.

We want to suggest that there is no tension between attunement and active inference. Our argument rests on the claim that active inference is what enables internal dynamics to attune to dynamics of the environment, and without implying the presence of internal representations (*contra* Helmholtzian anticipation). The AIF effectively takes the form of the principle of least action, where seeking to occupy local minima of surprise yields critical transition points that can be leveraged to facilitate occupation of expected states, and therefore to militate against undesired (aka surprising) states. Running either too fast or too slow when out to catch a fly ball induces surprise. Conversely, combating surprise is to chart a course of least resistance or action; namely, to settle on action policies enabling one to keep the fly ball stationary on the retina. Under the AIF, attunement in this specific sense is formulated in terms of a generalized descent on the free energy of internal states of a system (Friston, 2014).

Gradient descent optimization can be understood in terms of Bayesian belief parametrization. Bayesian belief parametrization or model optimization can be captured by the KL-divergence, $D_{\mathrm{KL}}\,[q(\mathrm{h}) \,||\, p(\mathrm{h|e})]$, which denotes the *relative entropy* between current beliefs or the prior density distribution, $q(\mathrm{h})$, and the (true) posterior distribution, $p(\mathrm{h|e})$. The KL-divergence is therefore a measure of the residual (or relative) surprise between the two probability distributions. When variational free energy is minimized, the prior density is approximately equal to the posterior distribution. The better this approximation, the smaller the divergence. This means that the prior density approximates exactly the same quantity that Bayesian inference seeks to optimize; namely, the posterior density distribution.

The AIF reminds us that it is not only internal (brain) states that engage in statistical inference, but also active states. This is what embeds the AIF in the context of both embodied and enactive views of mind, associating action with the process of selectively sampling sensory data to reduce surprise.

Some think that the KL-divergence underpins a representational view of statistical inference. The most articulated view of this representational reading comes via Kiefer and Hohwy (2018). Keifer and Hohwy state that any inference 'of the states of the world [or body] given by q(h) is correct when q(h) corresponds to what Bayesian inference would yield, i.e., when KL(q|h) p(h|e)) = 0' (*ibid.*, p. 20). From this they then argue that the KL-divergence provides a measure of misrepresentation. As they say: 'as long as KL(q|h) p(h|e)) > 0, the inferred state of the world given by q(h) is a misrepresentation' (*ibid.*, p. 21). It is therefore easy to see why Gallagher would think that optimization schemes like predictive processing and the AIF are committed to representationalism.

This conclusion, however, need not follow. To see this consider that the KL-divergence is the relative (Shannon) entropy between two distributions. Relative entropy is the divergence of the KL-divergence: q(h|e) and p(h|e), KL(q(h)||p(h|e)). Shannon entropy (or self-information) captures the insight that a variable is a source of information if it has a distribution of values, and that a variable can be said to express information about another variable in so far as the two variables are correlated (Godfrey-Smith and Sterelny, 2016). This is an important point, for it allows for the following observation: namely, that relative entropy can be spelled out as information in the form of covariance. The key idea of Shannon information is equivalent with long-term surprise. In the context of gradient descent optimization the prior distribution carries more information about the posterior distribution to the extent that it can be reliably used to approximate the posterior distribution conditioned on the prior and the posterior being systematically correlated. So the KL-divergence is an informational measure that equates to information as covariance, given that the tighter the divergence is the higher the covariance is between the correlated variables, and their associated values. This is a nice outcome, for it highlights that high covariance implies a shrinkage of surprise. Crucially, Shannon information does not have any representational or semantic implications. If it is information in the sense of Shannon that the KL-divergence is a measure of, then one cannot also say that one variable misrepresents or carries false information about another variable (*ibid.*; Kirchhoff and Robertson, 2018; Hutto and Myin, 2013). Hence, when Bayesian belief optimization is cast via the AIF, associating inference with inference over representational states need not follow (see Kirchhoff and Robertson, 2018, for further details).

We want to end by adding a further observation. One might insist that the higher one travels up the generative hierarchy of probabilistic inference, the more likely it is that there will be a need to associate inferences with representations. We do not wish to rule this point out here (for discussion, see Constant *et al.*, 2019). Instead we shall highlight two crucial points. The first is that it will not suffice to stick with a story about relative (Shannon) entropy to vindicate representational inference at the more abstract levels of probabilistic inference. This follows from our previous assessment that relative entropy can be expressed in terms of covariance, and covariance is not a free ticket en route to a representational meal (Kirchhoff and Robertson, 2018). The second is that the selective pruning of action policies can be associated with the context-sensitive formation of habitual action policies. A proposal is that habitual action policies are located at lower or shallower scales of hierarchical organization. These shallower parts of the hierarchy are much more specialized than operations at higher scales of organization, suggesting that far fewer selections are needed to select for relevant sequences of action. One way to think about this is that the AIF results in a kind of hierarchical embedding of non-representational action routines.

4. Conclusion

This paper argued that anticipatory action is inferential *qua* the active inference framework. Accepting this can illuminate a tenable way of addressing the acuity problem. The acuity problem, recall, is the puzzle of explaining that, on the one hand, fast, skilful action seems, when well-honed, to unfold automatically, reflexively, and without conscious thought or reflection, and yet, on the other hand, we seem often to characterize skilful actors by virtue of their capacities to anticipate and adjust to novelty or unexpected occurrences, often creatively and intelligently, under highly exigent circumstances.[12]

If the intelligently anticipatory dynamics involved in fast, skilful action approximate — on average and over time — a form of Bayesian inference, we can describe automatically-executed and highly habitual forms of action as intelligent in the sense that they emerge as the result of and continue to be instantiated by a process that adheres — to some extent — to rational (Bayesian) norms. Appeal to active inference allows us to explain the intelligence behind the fluid and context-sensitive dimension of skilful action, while, simultaneously, acknowledging that although the more habitual, reflexive responses on the part of the skilled actor will correspond to the same intelligent process, they have become highly engrained over time. Perhaps one would be inclined to characterize the agent as no longer engaged in inference when their action is highly habitual and occurs in a seemingly context-insensitive fashion (even if these context-insensitive responses have aided in an approximation to inference over the long term) (Orlandi, 2014). But the more interesting issue, we submit, is that active inference dictates that both the fluid, context-sensitive and the reflexive, unreflective dimensions of skilful action will both occur in approximate conformity with Bayesian norms. Skilful action is, then, inferential in a non-trivial (technical) sense.

We have argued that we should resist characterizing intelligent action as arising by virtue of participating in or facilitating Helmholtzian inference. The body plant is not enslaved in order to vindicate the hypotheses generated by the scientist-like brain. We have argued, *contra* Helmholtzian renditions of PP, that we should resist the notion of the brain executing an unconscious Bayesian inference over contentful hypotheses. We have argued, moreover, that characterizing anticipatory action as inferential *qua* active inference does not imply cognition to be realized entirely by neural dynamics. We conclude,

12 Often philosophers distinguish skilful from habitual forms of action on the basis that skills are intelligent in a way that habits are not (see e.g. Stanley and Krakauer, 2013). The AIF might go about showing habitual action to be intelligent in the sense that it will correspond to action policies that have, on average and over time, allowed the predictive organism to inferentially navigate their specific environmental niche. In this case, the seeming divide between automatic and intelligent action highlighted by the acuity problem might not seem as large. Even highly engrained habitual actions will qualify as rational in the sense that they have, by way of a process that conforms to (Bayesian) norms of rationality, aided the organism on average and over time.

then, that our inferential understanding of anticipatory action is substantially less intellectualist than competing accounts that understand the cognitive basis of anticipatory action as inferential.

Finally, and looking to the future, in this paper we did not provide an exhaustive mechanistic and information-theoretic account of the implementation of active inference. Rather, we have taken first steps towards addressing the acuity problem in terms of active inference. In future work we will augment this account by providing a formal and implementationally-informed framework for how the active inference paradigm not only addresses but solves the acuity problem.

Acknowledgments

Kirchhoff's work was supported by an Australian Research Council Discovery Project 'Mind in Skilled Performance' (DP170102987). Robertson's work was supported by an Australian Research Council PhD scholarship, as part of the Discovery Project 'Mind in Skilled Performance' (DP170102987). Thanks to Katsunori Miyahara and an anonymous reviewer for helpful comments.

References

Allen, M. (2018) The foundation: Mechanism, prediction, and falsification in Bayesian enactivism: Comment on 'Answering Schrödinger's question: A free-energy formulation' by Maxwell James Désormeau Ramstead et al., *Physics of Life Reviews*, **24**, pp. 17–20.

Allen, M., Levy, A., Parr, T. & Friston, K.J. (2019) In the body's eye: The computational anatomy of interoceptive inference, *BioRxiv*, 603928.

Anderson, M.L. (2014) *After Phrenology*, Cambridge, MA: MIT Press.

Bahill, T. & LaRitz, T. (1984) Why can't batters keep their eyes on the ball, *American Scientist*, **72** (3), pp. 249–253.

Bar, M. (2007) The proactive brain: Using analogies and associations to generate predictions, *Trends in Cognitive Sciences*, **11** (7), pp. 280–289.

Beer, R.D. (2000) Dynamical approaches to cognitive science, *Trends in Cognitive Sciences*, **4** (3), pp. 91–99.

Brooks, R.A. (1999) *Cambrian Intelligence: The Early History of the New AI*, Cambridge, MA: MIT press.

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** (6), pp. 2417–2444.

Burr, C. & Jones, M. (2016) The body as laboratory: Prediction-error minimization, embodiment, and representation, *Philosophical Psychology*, **29** (4), pp. 586–600.

Chapman, S. (1968) Catching a baseball, *American Journal of Physics*, **36** (10), pp. 868–870.

Chen, A.G., Benrimoh, D., Parr, T. & Friston, K. (2019) A Bayesian account of generalist and specialist formation under the active inference framework, *BiorXiv*, [Online], http://dx.doi.org/ 10.1101/644807.

Clark, A. (1997) *Being There*, Cambridge, MA: MIT Press.

Clark, A. (2013) Whatever next? Predictive brains, situated agents, and the future of cognitive science, *Behavioral and Brain Sciences*, **36** (3), pp. 181–204.

Clark, A. (2015) Predicting peace: The end of the representation wars, *Open MIND*, Frankfurt am Main: MIND Group.

Clark, A. (2016) *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*, Oxford: Oxford University Press.

Clark, A. (2017) Busting out: Predictive brains, embodied minds, and the puzzle of the evidentiary veil, *Noûs*, **51** (4), pp. 727–753.

Clark, A. & Grush, R. (1999) Towards a cognitive robotics, *Adaptive Behavior*, **7** (1), pp. 5–16.

Colombetti, G. (2014) *The Feeling Body: Affective Science Meets the Enactive Mind*, Cambridge, MA: MIT Press.

Colombo, M., Elkin, L. & Hartmann, S. (2018) Being realist about Bayes, and the predictive processing theory of mind, *The British Journal for the Philosophy of Science*, axy059.

Conant, R.C. & Ashby, W.R. (1970) Every good regulator of a system must be a model of that system, *International Journal of Systems Science*, **1** (2), pp. 89–97.

Constant, A., Ramstead, M., Veissière, S. & Friston, K. (2019) Regimes of expectations: An active inference model of social conformity and decision making, *Frontiers in Psychology*, **10**, art. 679.

Croft, J.L., Button, C. & Dicks, M. (2010) Visual strategies of sub-elite cricket batsmen in response to different ball velocities, *Human Movement Science*, **29** (5), pp. 751–763.

Di Paolo, E., Buhrmann, T. & Barandiaran, X. (2017) *Sensorimotor Life: An Enactive Proposal*, Oxford: Oxford University Press.

Dienes, Z. & McLeod, P. (1993) How to catch a cricket ball, *Perception*, **22**, pp. 1427–1439.

Dreyfus, H.L. (2002) Intelligence without representation — Merleau-Ponty's critique of mental representation: The relevance of phenomenology to scientific explanation, *Phenomenology and the Cognitive Sciences*, **1** (4), pp. 367–383.

Fink, P.W., Foo, P.S. & Warren, W.H. (2009) Catching fly balls in virtual reality: A critical test of the outfielder problem, *Journal of Vision*, **9** (13), art. 14.

Fridland, E. (2017) Automatically minded, *Synthese*, **194** (11), pp. 4337–4363.

Friston, K. (2010) The free-energy principle: A unified brain theory?, *Nature Reviews Neuroscience*, **11** (2), art. 127.

Friston, K. (2013) Life as we know it, *Journal of the Royal Society Interface*, **10** (86), 20130475.

Friston, K. (2012) A free energy principle for biological systems, *Entropy*, **14** (11), pp. 2100–2121.

Friston, K. (2017) The variational principles of action, in Laumond, J.-P., Mansard, N. & Lasserre, J.-B. (eds.) *Geometric and Numerical Foundations of Movements*, pp. 207–235, Cham: Springer.

Friston, K. (2018) Does predictive coding have a future?, *Nature Neuroscience*, **21** (8), art. 1019.

Friston, K., Adams, R., Perrinet, L. & Breakspear, M. (2012) Perceptions as hypotheses: Saccades as experiments, *Frontiers in Psychology*, **3**, art. 151.

Friston, K.J., Lin, M., Frith, C.D., Pezzulo, G., Hobson, J.A. & Ondobaka, S. (2017) Active inference, curiosity and insight, *Neural Computation*, **29** (10), pp. 2633–2683.

Froese, T. & Ikegami, T. (2013) The brain is not an isolated 'black box,' nor is its goal to become one, *Behavioral and Brain Sciences*, **36** (3), pp. 213–214.

Gallagher, S. (2017) *Enactivist Interventions: Rethinking the Mind*, New York: Oxford University Press.

Godfrey-Smith, P. & Sterelny, K. (2016) Biological information, in Zalta, E.N. (ed.) *The Stanford Encyclopedia of Philosophy*, [Online], https://plato.stanford.edu/archives/sum2016/entries/information-biological/.

Gibson, J.J. (1979/1986) *The Ecological Approach to Visual Perception*, Boston, MA: Houghton Mifflin.

Hayhoe, M.M., McKinney, T., Chajka, K. & Pelz, J.B. (2012) Predictive eye movements in natural vision, *Experimental Brain Research*, **217** (1), pp. 125–136.

Helmholtz, H. von (1867/1962) Concerning the perceptions in general, in *Treatise on Physiological Optics*, vol. III, 3rd ed., Southall, J. (trans.), New York: Dover.

Hohwy, J. (2012) Attention and conscious perception in the hypothesis testing brain, *Frontiers in Psychology*, **3**, art. 96.

Hohwy, J. (2013) *The Predictive Mind*, New York: Oxford University Press.

Hohwy, J. (2016) The self-evidencing brain, *Noûs*, **50** (2), pp. 259–285.

Hohwy, J. (2019) Quick'n'dirty or slow and rich? Clark on predictive processing and embodied cognition, in Irvine, L., Stapleton, M. & Colombo, M. (eds.) *Andy Clark and His Critics*, Oxford: Oxford University Press.

Hohwy, J. & Michael, J. (2017) Why should any body have a self?, in de Vignemont, F. & Alsmith, A. (eds.) *The Subject's Matter: Self-consciousness and the Body*, Cambridge, MA: MIT Press.

Hubbard, A.W. & Seng, C.N. (1954) Visual movements of batters, *Research Quarterly. American Association for Health, Physical Education and Recreation*, **25** (1), pp. 42–57.

Hurley, S. (2001) Perception and action: Alternative views, *Synthese*, **129** (1), pp. 3–40.

Hutto, D.D. (2018) Getting into predictive processing's great guessing game: Bootstrap heaven or hell?, *Synthese*, **195** (6), pp. 2445–2458.

Hutto, D.D. & Myin, E. (2013) *Radicalizing Enactivism: Basic Minds without Content*, Cambridge, MA: MIT Press.

Hutto, D.D. & Myin, E. (2017) *Evolving Enactivism: Basic Minds Meet Content*, Cambridge, MA: MIT Press.

Kiefer, A. (2017) Literal perceptual inference, in Metzinger, T. & Wiese, W. (eds.) *Philosophy and Predictive Processing: 17*, Frankfurt am Main: MIND Group.

Kiefer, A. & Hohwy, J. (2018) Content and misrepresentation in hierarchical generative models, *Synthese*, **195** (6), pp. 2387–2415.

Kirchhoff, M.D. (2015) Species of realization and the free energy principle, *Australasian Journal of Philosophy*, **93** (4), pp. 706–723.

Kirchhoff, M.D. (2017) Predictive processing, perceiving and imagining: Is to perceive to imagine, or something close to it?, *Philosophical Studies*, [Online], 10.1007/s11098-017-0891-8.

Kirchhoff, M.D. (2018) Hierarchical Markov blankets and adaptive active inference: Comment on 'Answering Schrödinger's question: A free-energy formulation' by Maxwell James Désormeau Ramstead et al., *Physics of Life Reviews*, **24**, pp. 27–28.

Kirchhoff, M.D., Parr, T., Palacios, E., Friston, K. & Kiverstein, J. (2018) The Markov blankets of life: Autonomy, active inference and the free energy principle, *Journal of the Royal Society Interface*, **15** (138), 20170792.

Kirchhoff, M.D. & Robertson, I. (2018) Enactivism and predictive processing: A non-representational view, *Philosophical Explorations*, **21** (2), pp. 264–281.

Kirchhoff, M.D. & Kiverstein, J. (2019a) How to demarcate the boundaries of the mind: A Markov blanket proposal, *Synthese*, [Online], 10.1007/s11229-019-02370-y.

Kirchhoff, M.D. & Kiverstein, J. (2019b) *Extended Consciousness and Predictive Processing: A Third-Wave View*, London: Routledge.

Knill, D.C. & Pouget, A. (2004) The Bayesian brain: The role of uncertainty in neural coding and computation, *Trends in Neurosciences*, **27** (12), pp. 712–719.

Land, M.F. & McLeod, P. (2000) From eye movements to actions: How batsmen hit the ball, *Nature Neuroscience*, **3** (12), art. 1340.

Linson, A., Clark, A., Ramamoorthy, S. & Friston, K. (2018) The active inference approach to ecological perception: General information dynamics for natural and artificial embodied cognition, *Frontiers in Robotics and AI*, **5**, art. 21.

Mann, D.L., Spratford, W. & Abernethy, B. (2013) The head tracks and gaze predicts: How the world's best batters hit a ball, *PLoS One*, **8** (3), e58289.

Mirza, B.M., Adams, R.A., Mathys, C.D. & Friston, K. (2016) Scene construction, visual foraging, active inference, *Frontiers in Computational Neuroscience*, **10**, art. 56.

Montero, B.G. (2016) *Thought in Action: Expertise and the Conscious Mind*, New York: Oxford University Press.

Orlandi, N. (2014) *The Innocent Eye: Why Vision is not a Cognitive Process*, Oxford: Oxford University Press.

Orlandi, N. (2016) Bayesian perception is ecological perception, *Philosophical Topics*, **44** (2), pp. 327–351.

Papineau, D. (2013) In the zone, *Royal Institute of Philosophy Supplement*, **73**, pp. 175–196.

Parr, T., Rees, G. & Friston, K.J. (2018) Computational neuropsychology and Bayesian inference, *Frontiers in Human Neuroscience*, **12**, art. 61.

Pezzulo, G., Kemere, C. & Van Der Meer, M.A. (2017) Internally generated hippocampal sequences as a vantage point to probe future-oriented cognition, *Annals of the New York Academy of Sciences*, **1396** (1), pp. 144–165.

Postma, D.B., Smith, J., Pepping, G.J., van Andel, S. & Zaal, F.T. (2017) When a fly ball is out of reach: Catchability judgments are not based on optical acceleration cancelation, *Frontiers in Psychology*, **8**, art. 535.

Rao, R.P. & Ballard, D.H. (1999) Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects, *Nature Neuroscience*, **2** (1), art. 79.

Ramstead, M., Badcock, P.B. & Friston, K. (2018) Answering Schrödinger's question: A free-energy formulation, *Physics of Life Reviews*, **24**, pp. 1–16.

Ramstead, M.J., Kirchhoff, M.D. & Friston, K.J. (2019) A tale of two densities: Active inference is enactive inference, *Adaptive Behavior*, 1059712319862774.

Ripoll, H. & Fleurance, P. (1988) What does keeping one's eye on the ball mean?, *Ergonomics*, **31** (11), pp. 1647–1654.

Seth, A.K. (2014) *The Cybernetic Bayesian Brain*, Frankfurt am Main: MIND Group.

Shaffer, D.M. & McBeath, M.K. (2005) Naive beliefs in baseball: Systematic distortion in perceived time of apex for fly balls, *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **31** (6), art. 1492.

Shapiro, L. & Spaulding, S. (2018) Embodied cognition and sport, in Cappuccio, M. (ed.) *Handbook of Embodied Cognition and Sport Psychology*, pp. 3–22, Cambridge, MA: MIT Press.

Stanley J. & Krakauer, J. (2013) Motor skill depends on knowledge of fact, *Frontiers in Human Neuroscience*, **7**, art. 503.

Stepp, N. & Turvey, M.T. (2010) On strong anticipation, *Cognitive Systems Research*, **11** (2), pp. 148–164.

Sutton, J., McIlwain, D., Christensen, W. & Geeves, A. (2011) Applying intelligence to the reflexes: Embodied skills and habits between Dreyfus and Descartes, *Journal of the British Society for Phenomenology*, **42** (1), pp. 78–103.

Thompson, E. (2007) *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*, Cambridge, MA: Harvard University Press.

Varela, F., Thompson, E. & Rosch, E. (1991) *The Embodied Mind*, Cambridge, MA: MIT Press.

Wolpert, D.M. & Flanagan, J.R. (2001) Motor prediction, *Current Biology*, **11** (18), pp. R729–R732.