

Bayes, Predictive Processing, and the Cognitive Architecture of Motor Control.

1. Introduction

Proponents of Bayesian theories of cognition argue that, at root, all perceptual and cognitive processes engage in Bayesian inference. *Hierarchical* views of the Bayesian mind suggest that the mind employs multiple levels of generative models to predict sensorimotor information that can be compared with feedback from sensory systems. Hierarchical versions of the Bayesian mind include both traditional generative Bayesian models (Tenenbaum et al., 2011), as well as the predictive processing implementation of Bayesian computation (Clark, 2015; Hohwy, 2013).¹

Bayesian theories have been particularly fruitful in explaining the functioning of the motor system. Behavioral success requires movements that are controlled in a fine-grained way, despite several sources of noise and uncertainty. On the theory of *optimal feedback control* (OFC; Todorov, 2006, 2009), the motor system employs generative models to predict the motor commands that will minimize a cost function. The predictive processing view puts forward a similar proposal, which, while not based on minimizing cost functions, posits a similar role for generative models in minimizing *surprisal* – a quantity measuring the deviation from predicted sensorimotor consequences that the organism experiences (Friston, Samothrakis, & Montague, 2012).

Aside from some discussion about whether modularity is compatible with the hierarchical approach (Drayson, 2017; Hipolito & Kirchhoff, 2019), or whether the predictive processing framework entails cognitive penetration (Lupyan & Clark, 2015; Vance & Stokes, 2017; Vetter & Newen, 2014), there has been little systematic discussion of the hierarchical picture's upshot for overall cognitive architecture.² On the standard construal of the hierarchical approach, it is tempting to see the mind as a continuous hierarchy, with each level of the hierarchy predicting activity at the one below it, out to the furthest extremities of perception and the motor system. Proponents of predictive processing differ on whether to embrace this conclusion – Clark (2015) seems to endorse it, while Hohwy (2013) thinks

¹ I will talk as generally as possible about frameworks based on hierarchical generative models, so I will aim to include both Bayesian and predictive processing approaches in the discussion. Technically, predictive processing is one way of *implementing* Bayesian inference, which is often computationally intractable in its pure form. Importantly, proponents of the PP framework take motor control generally, and OFC in particular, to be compatible with their approach. So, I will discuss them in tandem, despite the fact that the PP framework differs from a purely Bayesian theory the formal details.

² Much critical discussion of the framework assesses its explanatory value in explaining cognition. While these issues are important, I will not engage them closely here; see (Cao, 2020; Colombo, Elkin, & Hartmann, 2018, 2020; Colombo & Hartmann, 2017; Colombo & Series, 2012; Kaplan & Hewitson, 2020; Zednik & Jäkel, 2016).

there is empirical and theoretical reason to posit that, at least at some points, distinct levels of the hierarchy are relatively encapsulated from each other.

The question of whether there is a continuous hierarchy is important, both internally to formulation of the Bayesian approach, and for how the view contributes to other questions. For instance, theorists in action theory have recently reached for Bayesian ideas as a way of potentially solving the “interface problem” – the problem of how abstract propositional intentions at the top of a representational hierarchy can cause specific motor behaviors at a much more fine-grained level. These proposed solutions in fact *rely on* the idea of a continuous processing hierarchy, and assume that the hierarchical generative model approach offers just such a proposal.

I have two related aims in this paper. The first, lower-order question I will assess is the question of whether there is a single, continuous hierarchy of processing in the mind, specifically involved in motor control. My answer to this will be ‘no’. On an alternative view of motor control, the motor system does not require input from distinct, higher-level cognitive representations or models to make decisions and implement motor plans. Instead, the motor system itself has the ability to select goals and initiate appropriate behavior. I will argue that this perspective better matches the literature on OFC, and thus that we should resist the continuous hierarchy approach. Hence, a broad embrace of hierarchical generative accounts should not be taken as a commitment to the continuous hierarchy view.

I also have a series of broader aims, however, which might be described as meta-level analysis of trends in theorizing about cognitive architecture. Theories based on hierarchical generative models, and in particular predictive processing-based versions of those theories, are often touted for their *unifying* power. They are taken as accounts of *all* perception, cognition, and action. Part of what I intend to show is that this unifying perspective is, in some senses, misleading. Embracing these views simply underdetermines many key questions we might want to ask about cognitive architecture, such as whether there is a continuous processing hierarchy. As such, caution must be used when applying insights from these views in other arenas, for instance in action theory.

I proceed as follows. In section 2, I discuss the background. I introduce (section 2.a) the hierarchical Bayesian approach, closely following its development within the predictive processing framework. I discuss how the “transitivity” of processing on views based on generative models might be taken to suggest the presence of a continuous hierarchy. I then (section 2.b) discuss the interface problem, and how philosophers of action have been motivated to see the continuous hierarchy as a solution to it. In section 2.c I further describe ways in which higher- and lower-level models might be distinguished, including distinctions in terms of the *form* of the models. In section 3 I introduce OFC, and in section

4 I describe the results that I take to show that sensorimotor inference works based on models with a sensorimotor format. In section 5, I discuss alternative perspectives on the relationship between cognitive and sensorimotor systems. Section 6 concludes.

2. The Background

2.a. Hierarchical Generative Model-Based Views of Cognition

In a cognitive architecture based on hierarchical generative models, the key notion is that, at each level of processing, a generative (or “forward”) model at one level predicts outcomes at the level(s) below it. The job of the system is to update its overall model, which is distributed across levels of the hierarchy, on the basis of *error signals* it receives bottom-up from the sensory system. The process of updating leads the system’s model to reflect the causal structure of the world more accurately, and to better predict sensory consequences on the basis of so doing.

The hierarchical picture gives rise to a kind of *transitivity* (cf. Drayson, 2017). If level A predicts the outputs of level B, and level B predicts the outputs of level C, then level A has, to some degree, predicted the outputs of C. If this process occurs all up and down the hierarchy, then it seems there is a continuous stream of prediction and error processing, beginning at the highest levels and ending at the very periphery of sensory and motor systems. Indeed, it is hard to know how to *avoid* this conclusion, on the view that this is just what processing *is*. Every level, on this view, needs to receive predictions and send back prediction error, so for any active level must be receiving input from one above.³

As one might expect on this view, many have suggested that there are no fundamental distinctions or gaps anywhere along this hierarchy – that distinctions such as “perceptual” versus “conceptual,” or “perception” versus “action,” are simply glosses on a more fundamental continuity of hierarchical inference (Clark, 2015; Vance & Stokes, 2017). States such as “beliefs” and “desires,” or “concepts,” are not distinct mental kinds from perceptual and motor representations, but only regularities represented in a hierarchical model at a higher level (cf. Anderson, 2017; Dewhurst, 2017).

Not everyone wants to come to that conclusion, however. Hohwy (2013) suggests that there are both empirical and functional reasons to posit limitations to the continuity of prediction. First, cases of impenetrable perceptual illusions suggest that, at least in some cases, higher-level knowledge fails to override sensory information. Second, functionally-

³ One worry that transitivity raises is a worry about *regress* – it seems that, if processing is to stop somewhere, it must be at a final arbiter who makes the final say, a kind of predictive homunculus. Hohwy and Clark attempt to avoid this problem by arguing that it is the hierarchical model as a whole that is met with sensory data. The regress worry is not my concern here, so I will not analyze these proposals in detail.

speaking, the system is better-off if some sources of information can be treated as independent sources of evidence, which a fully continuous hierarchy seems to preclude. Hohwy also offers the helpful thought that sometimes quite general assumptions might be embodied by perceptual systems themselves; he gives as an example the “light comes from above” prediction that explains inverted-face perception and other effects. If this is true, then the perceptual system does not necessarily require general assumptions to be sent top-down from more “cognitive” systems.

If one thinks that some parts of the mind are, at least relatively, encapsulated from other parts, then one needs to posit some kind of break in the transitivity of hierarchical prediction.⁴ But it is not clear how to characterize this break. The only suggestion I can find in the literature is one based on differences in *grain* between distinct levels. A common suggestion in the Bayesian/predictive processing literature is that higher levels of the hierarchy encode regularities (and concomitant predictions) at higher levels of spatiotemporal grain than lower levels. At the lowest levels, the system has sensory areas making predictions about a particular kind of sensory stimulation within tiny receptive fields. At the highest levels, you have deep convictions about the broad structure of the physical, perceptual, and social world. Hohwy (2013) suggests that, in some instances, higher-level predictions simply can’t predict sensory consequences fine-grained enough to overcome bottom-up signal. This explains, among other things, the inability to overcome perceptual illusions by understanding how they work.

Unfortunately, as far as I can tell, no one has attempted a systematic discussion of when levels of grain might fail to predict lower-level input, and it is not obvious under what circumstances one would expect failures in transitivity due to difference in grain. Given that each level needs a predictive input from above, this assumes that each level can, to some degree, either directly or indirectly, reach the levels below it. Put another way, if granularity is a problem in one instance, then why is it not a problem in *all* instances? Each higher level, after all, is supposed to encode regularities at a broader spatiotemporal grain. How can they do this *and* predict the details at the lower level? So, it seems like either each level can predict lower levels, and transitivity takes hold (along with its suggestion of a continuous hierarchy), or they cannot, and the entire picture has a serious problem.

So, it is not obvious whether one should construe the hierarchical generative framework as entailing a continuous processing hierarchy. In the next section, however, I will show how

⁴ I say “relatively encapsulated” here because the kind of encapsulation I am discussing is, I believe, weaker than modularity traditionally construed. It is possible for there to be relatively important causal connections between two mental systems, where that influence is not best described as their computing over each other’s contents. Indeed, in other work I describe in detail how this works in the relationship between cognition and perceptual/motor systems ([author’s papers]), and I further describe this kind of relationship below.

some theorists interested in the psychology of action have assumed that just such a picture follows from the framework.

2.b. Hierarchies and Intentional Action.

At first glance, considerations in the philosophy of action may seem far removed from the details of cognitive architecture. However, a recent trend in philosophy of psychology has attempted to legitimate views from the philosophy of action – in particular, Davidson’s (1980) *causal* theory of action – with reference to psychological theory. On the causal theory, what individuates intentional actions is their causal history. An action is intentional if and only if it is caused, in the right way, by the agent’s attitudes. These attitudes are the standard folk psychological list, including beliefs, desires, and, particularly, intentions. So, explanation of actions cites the agent’s reasons for acting, and their reasons cause their actions through a causal interaction between their attitudes and their body.

How would such a picture be realized in a cognitive architecture? One popular proposal is to make a distinction between *motor* and *practical reasoning* systems (Pacherie, 2008). The practical reasoning system is in charge of deliberation, and operates via syllogistic reasoning. A *decision* happens in this system, and the result of that decision is an intention to act. The problem of action is then reduced to describing how the content of the intention is propagated to the motor system. The intention must cause the motor system to act in a particular way, so as to carry it out. Mylopoulos and Pacherie (2017, 2018) describe this as the need for a “content-preserving causal process” relating the intention and the motor system command that carries the intention out.

One worry for the picture above is how the “content preserving causal process” works. On this view, the decision needs to be propagated to the motor system via an intention, but depending on how you type the states involved, this can appear problematic. For instance, one posited role for intentions is that they are *abstract* and range over particular circumstances – my intention to eat ice cream is supposed to be the same across different instances of eating ice cream, even though the particular circumstances may differ. But how does the abstract intention control and bring about bodily movements that are, inherently, tied to particular circumstances? I have to move just such-and-such a way to grasp a particular ice cream cone, but the abstraction of the intention seems to prevent it from describing this particular circumstance. This has come to be known as the “interface” problem (Burnston, 2017; Butterfill & Sinigaglia, 2014; Ferretti & Caiani, 2018; Mylopoulos & Pacherie, 2017).

Where this intersects with the hierarchical generative framework is that, in the interface debate, theorists are precisely *trying to establish* the continuity of that hierarchy. One common way of solving the interface problem is to posit a series of intervening stages of

processing that mediate the transition from the practical reasoning system to the motor system. And, unsurprisingly, theorists have sometimes reached for Bayesian ideas in constructing these views. For instance, Mylopoulos and Pacherie suggest that a process of hierarchical Bayesian learning eventuates in the coordination, via a range of “action concepts” at the interface of the practical reasoning and motor systems. Pavese (2017) cites the prevalence of forward (generative) models in theories of motor control to argue that intentions are translated into basic motor commands.

It is worth noting that those primarily interesting in the Bayesian approach have come to similar conclusions. Rescorla (2016), for instance, thinks that discussion of propositional attitudes is *endemic* to Bayesian theorizing about motor control. Motor control, on Bayesian views, is always defined relative to a *goal* representation. It is only relative to a particular goal that one can construct a cost function and infer the optimal action to minimize it. But, Rescorla correctly points out, the models often take the goal as given, or constrain it highly with the task setting. They often do not say where the goal representation comes from. Rescorla thinks the only place it can possibly come from is a prior decision on the part of the actor. He says, “intention plays a particularly central role within optimal control modeling. The main idea behind the optimal control framework is that motor activity chooses commands suited to advance some goal. This idea underlies the framework’s distinctive explanations, including the explanation for task-constrained variability. In many cases (though perhaps not all), the relevant goal is set by higher-order cognition.” (Rescorla, 2016, p. 19).

So, we have two distinct philosophical discussions revolving around the same issue – to what extent is there a continuous hierarchy of probabilistic representation spanning the “highest” levels of cognition and the “lowest” levels of sensorimotor representation? If one wants to solve the interface problem, then one is encouraged by frameworks that suggest a continuous hierarchy. If one is a Bayesian but believes in the virtues of (perhaps partial) informational encapsulation at lower levels, one might be inspired to resist the continuous hierarchy. I think the continuous hierarchy view should be resisted. In the remainder of the paper, I give an argument for this view, and conclude by suggesting what the effect denying the continuous hierarchy has on the surrounding discussions.

2.c. Representation and the Bayesian Framework.

One of the major ongoing debates in the Bayesian brain literature focuses on whether the framework is best understood as positing mental representations, or not. Those who argue that it is representational focus on the notion of “message passing” between levels as implying a representational view (Hohwy, 2013). Further, they have argued that notions such as misrepresentation can be understood within the predictive processing framework (Kiefer & Hohwy, 2017), and that a notion of content can be given based on the idea of

structural resemblance between hierarchical models and the world (Kiefer & Hohwy, 2019; Gladziejewski & Milkowski, 2017; cf. Williams, 2018b).

Anti-representationalists about the framework object to the “internalist” aspects of the representational construal of hierarchical models (Kirchhoff, 2018; Kirchhoff & Robertson, 2018). They argue that this reading implies that the mind is ensconced from the environment, and hence that it fails to capture the ways in which hierarchical generative systems are embedded in a world and active participants in it. These theorists also deny that the necessity of a representational view of inference, and that formal notions of discrepancy between a model and the world amount to misrepresentation. Rather than viewing hierarchical models as the internalist does, they suggest, we should construe organisms’ bodies as partially constitutive of their cognitive models, thus denying the need for an organism’s internal models to represent their bodies and their surroundings (Friston, 2011; Friston, Thornton, & Clark, 2012; Kirchhoff & Robertson, 2018).

I wish to remain as neutral as possible in the debate, but I will occasionally resort to representation-talk when convenient, so I want to clarify here how I am construing the notion. I am committed to the following: (i) hierarchical models are structurally and functionally decomposable; (ii) their functions are individuated partially by what things they predict about the body and/or the environment, and how; (iii) causal relations between distinct levels of models are part of the hierarchical generative model-based account of how behavior arises.

I think (i)-(iii) are commitments across the board – non-representationalists about PP, for instance, still endorse the hierarchical-model based account, and they do not deny that levels in these models differ partially in the kinds or levels of generalizations they predict. The question is whether one thinks (i)-(iii) *must* be interpreted representationally. My preferred view is that there is a less substantive notion of representation than is traditionally embraced in cognitive science, which plays an explanatory role here (cf. Anderson’s control” notion of content; Anderson, 2017), and which can incorporate some of the important insights from “externalist” PP, but I don’t intend to defend that position. Rather, I will talk about levels of models as employed in optimal feedback control, and assume that either a representational or a non-representational view will have to make sense of these models within the aegis of commitments (i)-(iii).

So, how are models at different levels construed? One way, discussed in section 2a, is via the notion of *grain*. I raised some problems for this view above. Another way, which is popular in the interface problem literature, is to individuate processes at distinct levels by their functional or representational *form*. In representational terms, one can construe this as a distinction between *modal* and *amodal* representations. On the form distinction, amodal representations are distinguished by referring in a different way than modal

representations. The structure of amodal representations is arbitrary with respect to what they represent. Conversely, sensorimotor representations involve *metrical* relationships to what they represent (cf. Author's paper). That is, they represent via placing the represented object as instantiating particular values along continuous dimensions that comprise the representation. Dimensions constitute a *space*—each represented object exists at a particular location in the space, determined by its values along the represented dimensions.

It is unclear exactly how non-representationalist PP would account for a distinction in format between distinct models, but I think the best approach would be along the lines of commitment (ii). Different models, or aspects thereof, might gain their distinct functional roles by differing in what aspects of the body or world they predict. That is, a metrical model might embody a distribution over specific dimensional values of a stimulus or motor action, whereas an amodal model would embody a distribution over category membership.

Regardless of how precisely they are explained, there is strong evidence that the broader Bayesian perspective is at least compatible with, and sometimes employs, distinctly structured models. Theorists within the hierarchical generative model-based approach often note this explicitly. Sprevak (2019), for instance, says that “in principle, probabilistic representations could use any physical vehicle, and any formal format” (p. 601). Clark (2015) also suggests that, in predictive processing, “there is no need to fix on any single form of knowledge representation. Instead, each layer is free to use whatever form of representation best enables it to predict and (thus) account for the activity at the level below” (p. 173). Similarly, while Tenenbaum and colleagues sometimes describe models of inferences using multiple nested abstract relations as employing a “probabilistic language of thought” (Goodman, Tenenbaum, & Gersteburg, 2014), other models employ distributions over detailed motor patterns (Lake et al., 2017). While not put explicitly in terms of the form distinction, these distinct kinds of models implicitly invoke differently structured models.

A major claim of this paper is that motor control is best understood as operating within metrically structured models. In the motor system, the quantities modeled include activation parameters for the individual muscles, orientation parameters for the limbs, and angle and torque parameters for the joints (Todorov, 2009). So, in the motor realm, ‘grasping’ might involve particular values for the orientation of the hand/fingers, the angle at the elbow, and the force of the grip. Generally, in action control, the relevant space is “sensorimotor” comprising both perceptual and motor spaces. This is because any goal-oriented action, such as a grasp, must be oriented towards a particular target, meaning that the muscular and joint parameters must be manipulated appropriately for a target defined along (e.g.) depth, shape, and orientation dimensions. I will only focus on “visuomotor” spaces in this paper, but other modalities could of course be involved.

I will explore the notion of a visuomotor space in more detail throughout the paper. For now, let us note that the interface problem is made more serious if one distinguishes propositional from sensorimotor representations/models by their form. On this view, practical reasoning is a process that takes place over propositionally structured representations/models, with amodal concepts as their constituents, while motor control happens in a distinct, sensorimotor format. But how, given the difference in format, can one representation cause detailed results in the other?

A variety of proposals have been made, but I have argued at length elsewhere ([Author's paper]) that these approaches fail to solve the "diversity/specificity" problem. That is, for any given conceptual state, say 'grasp', there are many detailed instantiations that might be implemented in a motor representation. Yet, one specific one out of those representations *must* be tokened for each grasping instance. So, there is no way to translate specifically from one type of model into the other. I will not rehash the argument here, other than to note that the interface problem is only a problem if one thinks that decisions and abstraction are proprietary to the practical reasoning system. If one denies this assumption, then there are a range of possible ways to think about the relationship between modal and amodal representations. In section 5, I will explore some of these. For now, I want to ask what it would look like to deny the assumptions just mentioned.

There are two aspects that are key to the alternative proposals I will consider. The first is a denial of the claim that *only* practical reasoning, and not the motor system, is involved in decision-making. There are several reasons to deny this assumption in the case of action control. For instance, Todorov (2009) suggests that action-control is *closed-loop*, rather than *open-loop* control. In open loop control, an entire trajectory of movement is pre-specified in advance, and the goal of the control system is simply to not deviate from that trajectory. However, this kind of system is bad at dealing with in-course adjustments, of which our motor systems are highly capable. Closed-loop control does not face this problem – it only specifies the end state, and the trajectory is continuously calculated at each time-step. If the motor system is closed loop, then figuring out what to do at each step is precisely what the motor system is in charge of. This is just to say that its operation is partially constitutive of the decision to perform a certain action in a certain way.

In concert with this perspective, Selen et al. (Selen, Shadlen, & Wolpert, 2012) have shown that activations in the motor system track the amount of *evidence* in favor of a certain action option, and hence that the motor system is sensitive to evolving evidence before an action target is determined. Moreover, it is physiologically well established that the motor cortex represents actions directed towards particular outcomes prior to the making of a decision, and that these representations are modified by how *valuable* the action's associated outcome is expected to be (Cisek & Kalaska, 2005). This kind of perspective suggests that decision-making is not the sole purview of a particular part of the brain, but instead is distributed

across several systems ([Author's paper]; Cisek & Kalaska, 2010). If this is possible, then perhaps the interface problem does not need to be solved at all ([Author's paper]).

The second aspect is to deny that *abstraction* is proprietary to a certain format of mental representation. While it is true that metric representations must represent particular values along their dimensions, that does not mean that no notion of abstraction is possible within a representational space. In particular, abstraction can occur in a space by (i) representing a *range* of values, or (ii) via *dimensionality reduction*. As we will see, motor control frequently operates by representing *variance* along a certain dimension or dimensions. But variance, and other statistical properties such as a mean, can be represented as probability distributions *across areas of representational space* (Körding & Wolpert, 2004; Zhang, Daw, & Maloney, 2015). And indeed, the possibility of representing probability distributions within visuomotor space is key to understanding the Bayesian perspective. Moreover, one can define *relations* between values along distinct dimensions, and doing so can lead to dimensionality reduction, because one variable will now suffice to describe the state of the system, vis-à-vis those relations.

Here is an example. A popular concept in the literature is that of a “motor synergy” (Bruton & O’Dwyer, 2018; d’Avella, Saltiel, & Bizzi, 2003; Overduin, d’Avella, Roh, Carmena, & Bizzi, 2015; Todorov & Ghahramani, 2004). Synergies are correlated movements of multiple muscles, which can be represented as a single variable – a kind of “control knob” that can be tweaked to govern multiple muscles in a coordinated fashion. The important point to note is that this correlation is itself defined in the motor space. So, changing values in the control knob is, by definition, twiddling the correlated dimensions that it relates. In constructing a synergy, one has represented a subspace of the overall space with a single variable, but that does not mean that the representation is operating outside of the space.

If abstraction is possible within a sensorimotor format of representation, this shows that the form distinction is different than the fine-ness of grain distinction. On the form distinction, distinct levels of grain might be represented within each of distinct systems (I discuss this further in section 5.) But that does not mean that they are functionally equivalent. Different formats of representation might be better or worse at different tasks, and therefore might be involve in generating behavior in different ways. In the remainder of the paper, I will be arguing that the best explanation for the range of effects within a broadly Bayesian approach to motor control posits that motor control operates within representation of sensorimotor space. This will open up new possibilities for thinking what kind of interactions amodal representations might have with this basic kind of operation. In the next section, I introduce the OFC approach to motor control in more detail.

3. OFC.

Todorov (2009) asks us to construe the “inference” in motor control the following way: if we assume that the goal has been attained, what is the sequence of bodily movements that led to that occurrence? So, given the sensory outcome of having attained the goal (e.g., with your fingers successfully wrapped around your coffee cup), what sequence of motor commands is likely to have generated those sensory consequences? This inference problem faces the issue of motor “redundancy” -- the fact that many different motor command sequences could eventuate in the same outcome. Think of the number of different arm angles from which you could begin your action of grasping your coffee cup. Each of these is a possible causal history for the sensory consequence of your hand grasping the cup.

Given this difficult problem, a popular proposal is that the motor system operates according to “optimal feedback control,” (OFC) which is the idea that motor inferences are constrained by a *cost function*. Usually, the cost function is construed as a combination of two quantities – an energy cost that penalizes inefficiency of movement, and a cost for deviation from the target (i.e., for decreased accuracy). The movement that is inferred is the one that minimizes the cost function. To calculate this, the motor system relies on a forward model of the dynamics and kinematics of the body, and the dynamics of the environment, which allows for the sensory consequences of distinct possible movements to be computed. Interestingly, it can be shown that the quantity that minimizes the cost function is equivalent to the posterior in a Bayesian inference (Todorov, 2009). That is, by showing what motor sequence best minimized the cost function, one has computed the most likely motor sequence for the generated outcome. Over time and repeated iterations, the forward model can be adjusted to more accurately predict the motor sequences that minimize the cost function. The kind of inference being performed here is compatible with, and is often explicitly cited as falling under the aegis of, predictive processing theories (Clark, 2015; Hohwy, 2013).

The process is hierarchical, in three senses. First, the use of a generative forward model guarantees that there must be an interaction between predicted consequences and sensory feedback. Second, the forward model operates across levels of spatiotemporal grain. That is, predictions are needed for both the immediate sensory consequences of a particular motor command, and the longer-term consequences of a series of motor commands. Third, the representations at higher levels are taken to be more “abstract,” in some sense, than the ones at lower levels. This is phrased a variety of ways, but one important notion is that the representations at higher levels generalize over variations in the representations at the lower levels.

So, how does the notion of hierarchical motor control fit into the considerations from the previous section, namely, the issue of whether there is a continuous hierarchy? There are two considerations within OFC that seem to speak in favor of the continuous hierarchy

view. The first is simply that there is a hierarchy graded in terms of levels of abstraction. Given that actions can be contemplated at any level of spatiotemporal grain (e.g., grasping the cup in front of you versus making it to the finish line of the marathon), there does not seem to be any “upper bound” of abstraction at which action control stops. Second, as Rescorla (2016) notes, all of the computations discussed in this section rely, necessarily, on a representation of the action goal or target. Without a description of the goal in mind, there is no end state against which to compare the inferred motor commands, relative to the cost function. Representation of the target goal is required to get the process running. At the intersection of these two points, one often sees the higher levels of the hierarchy described in more abstract terms, as employing “rules” or “symbols” (Braun, Mehring, & Wolpert, 2010; Haruno, Wolpert, & Kawato, 2003), which are more abstract than the motor primitives that operate at lower levels. If one is inclined towards the continuous hierarchy view, one would argue that these considerations suggest a stream of processing that begins at person-level, propositional states, and culminates at computations of basic motor actions.

However, there are also several aspects of the OFC framework that are not obviously in keeping with the continuous hierarchy view. In particular, many results suggest that computations within the OFC framework take place in a sensorimotor space. Importantly, the *action targets* and *motor commands* are represented in this space. Commands are often construed as transitions between motor parameters and sensory consequences. And this is important for how the computations work – energy costs are represented in terms of the motor parameters that lead to them; accuracy costs are computed by deviation from the particular sensory consequences that constitute successful performance of the action. Similarly, Bayesian approaches are tied to representations of *uncertainty* in the forward model; that is, in order to perform Bayesian inference the system must represent its degree of uncertainty in the outcomes of its movements and in the sensory data. But, Zhang et al. (2015) argue that uncertainty is represented as a probability distribution defined over “possible movement outcomes in space or time” (p. 1152), again suggesting that the computation happens in the sensorimotor space (cf. Merel, Botvinick, & Wayne, 2019).

The upshot of this is as follows: if goal states and abstraction are represented in the sensorimotor space, there is no guarantee that those representations are sensitive to processing at all levels of the hierarchy. The representation of a particular set of sensory consequences in a sensorimotor space is not, I take it, what philosophers normally mean by “intention.”⁵ And it is worth noting that I cannot find any substantive uses of the term

⁵ One can, of course, loosen the notion of ‘intention’ to allow it to encompass this view (Rescorla, 2016, for instance, suggests that he is neutral with regards to what kind of state an intention is). The problem with loosening the individuation conditions on intentions in this way is that the notion may now range over distinct kinds of mental states with distinct functional roles, thus preventing ‘intention’ from being useful for specifying particular causal/functional states in the cognitive architecture of action control. For further discussion, see ([Author’s paper]).

'intention' in any of Todorov's or Wolpert's writings. That doesn't mean that the hierarchy does *not* work on intentions, but positing that it does is an addition put in by philosophers that has no obvious formal corollary in the models, which involve transitions between sensorimotor states towards a perceptually represented goal.

In what remains of the paper, I want to argue that motor processing is in fact highly *discontinuous* with thought based on concepts and propositional attitudes, and hence that we should resist the continuous hierarchy view. The argument will be based on the idea that most of what motor processing computes is computed within sensorimotor space, and therefore that (i) the motor system does not *need* higher-level predictions for most of what it does, and (ii) that there are reasons to think that higher cognitive representations have limited *access* to the workings of the motor system in that space. In the next section, I discuss several powerful experimental paradigms in investigation of visuomotor control, and suggest that each shows that forward model-based motor control happens within the visuomotor space. I will then pursue what this question means for larger questions in cognitive architecture.

4. Some Experiments.

Experimental paradigms within a broadly Bayesian approach to motor control consist in manipulations meant to isolate the forward model at work, and to see how it changes with experience. Subjects are asked to perform motor movements, for instance moving a joystick, or other "manipulandum" so as to navigate a computer cursor towards a target. A range of manipulations are then possible. One can vary visual feedback by changing the position of the cursor mid-trial. Or one can change the properties of the manipulandum, subjecting it to various forces. In turn, fine-grained measurements can be taken, including the trajectories of the movements, the force exerted on the joystick, or (via EMG) the activation of the muscles. The goal is to see, by subjecting subjects to a range of manipulations, what the properties of their forward models are and how those models change in the course of learning tasks. In this section, I discuss three types of experiments, focusing on perturbations, generalization, and context. (I will discuss them in turn, although the relevant methods often overlap within an individual study.) I will argue that the best explanation for these results is that motor control inferences are performed in representations of visuomotor space.

4.1. Perturbation

Franklin and Wolpert (2008) had subjects move a manipulandum to guide a visually presented computer cursor towards a target. They then introduced a perturbation to the visual feedback of the hand position, i.e., they shifted the cursor during a trial. This varied in three ways. The perturbation was either to the left or the right, it was either task-relevant

or task-irrelevant, and it was either early or late in the trial. (Task relevance was determined by whether the perturbation persisted through the trial, thus requiring a response, or resolved by itself, hence not requiring one.) By measuring EMG signals of the muscles and force exerted on the manipulandum, Franklin and Wolpert measured how subjects reacted to the perturbations.

Subjects showed corrective responses to the perturbation *only* in the task-relevant condition, and their responses were directionally specific to account for the direction of perturbation. Moreover, the responses were adjusted depending on whether the perturbation occurred early or late in the trial, and hence whether a greater or lesser force was required to correct the movement by the time the target was reached. Importantly, though, EMG signals specific to the adjustment emerged as quickly as 90ms after the perturbation, showing that the responses occurred *prior* to conscious recognition and voluntary adjustment on the part of the subject. On this view, the motor system itself keeps track of task relevant variation and modulates behavior accordingly.

This kind of result can be expanded in a number of ways. Diamond et al. (Diamond, Nashed, Johansson, Wolpert, & Flanagan, 2015) have shown that such involuntary adjustments to perturbations depend on *object dynamics*. They created novel “objects” by changing the force loads on the manipulandum, and showed that, after a short learning period, subjects responded to visual feedback perturbations in a way that accounted for these dynamics. In particular, the authors suggest that the adjustments are made in a way that would maintain grasp stability if the object were freestanding. Hayashi et al. (Hayashi, Yokoi, Hirashima, & Nozaki, 2016) replicated the finding of involuntary, task-relative adjustments with different kinds of perturbations (e.g., rotations) to the visual feedback.

These effects show that fine-grained adjustments to motor behavior (i) track task-relevant variation in the occurrent environment, and (ii) adjust for this variation on the basis of a forward model. But since the *adjustments* that the motor system makes in real time are in visuomotor space, the best explanation for the effects is that the forward model is itself a representation of visuomotor space. And this is indeed how they are described in the literature.

4.2. Generalization.

Abstraction is importantly related to *generalization* – a representation is more abstract if it ranges over particular instances. In generalization experiments, subjects are presented with a range of conditions, and if subsequent performance in new, similar conditions is facilitated, one concludes that the subject’s forward model has been abstracted from the range of training examples. Wolpert and colleagues (Braun et al., 2010; Braun, Waldert, Aertsen, Wolpert, & Mehring, 2010) attribute this phenomenon to *structural learning*, in

which subjects learn a lower-dimensional structure that captures variation across specific instances. In this subsection, I argue that these abstractions are themselves represented in sensorimotor space.

In an early study of generalization, Körding and Wolpert (2004; cf. Hewiston et al., 2018) introduced a range of feedback perturbations to subjects as they performed a motor task, which required a motor response. After training, subjects in particular trials would adjust towards the *mean* of those perturbations, and their adjustments would be closer to the mean when the sensory feedback was uncertain (e.g., presented with visual noise or blur). The conclusion from this is that subjects' forward models represent the statistical properties of the perturbation and cause subjects to adjust accordingly, taking uncertainty into account. But the distribution of perturbations, and hence the mean, are precisely represented in the visuomotor space – it is their locations in this space across trials that determines their statistical properties. Thus, this kind of learning and generalization happens within the representation of the visuomotor space.

This kind of paradigm can be expanded in a number of ways. Braun et al. (Braun, Aertsen, Wolpert, & Mehring, 2009) subjected subjects to a range of perturbations in visual feedback during a reaching task, namely “rotations” of the cursor at varying degrees. After training, this experimental group, as well as a control group who had not received varying feedback during training, were exposed to a perturbation at a novel degree rotation. Subjects in the experimental group were both faster and more accurate than those in the control group. Braun et al. explain this as a general task structure (rotations in visuomotor space) applying across training conditions, with only a parameter (actual angle of rotation) changing. Since subjects in the experimental group learned this structure, they were able to perform better on a new example.

Further, this kind of result applied to *exploration* in novel contexts. In another experiment, Braun et al. compared two groups who underwent different training regimens. One group was subjected to a range of horizontal perturbations in the cursor position, and one to a range of vertical perturbations. After training, each group was then given novel trials for either the structure they *had* learned, or the structure they *hadn't* learned. For a perturbation within their previously learned structure, subjects were faster and exhibited less variance in their trajectories towards the target. For a permutation in the not-previously-learned structure, not only were subjects less accurate, with more variance in their trajectories, but the added variance tended to occur *withing the previously learned* structure. That is, subjects trained in horizontal perturbations would exhibit variance in the horizontal structure when moving towards a new, vertically-perturbed target, and vice versa. Braun et al. take this as evidence that, when faced with a new scenario, subjects preferentially explore novel trajectories according to the structure they have previously

learned. This preferential search along learned structure has also been observed in a variety of visuomotor and attention tasks (for further discussion see [Author's paper]).

Further investigation of generalization has proceeded by analyzing the relationship between generalization in *intrinsic*, or bodily, space, and *extrinsic*, or visual space. These can be varied relative to each other by keeping the body position the same while varying visual stimuli, or by doing the reverse. Some results suggest that, for certain tasks, generalization can occur entirely within a representation of extrinsic space, for instance by showing that learned forward models generalize across limbs (Hewitson, Sowman, & Kaplan, 2018). For other kinds of tasks, generalization is strongest only when novel tasks are similar in both intrinsic and extrinsic space (Carroll, Poh, & de Rugy, 2014). Brayanov et al. (Brayanov, Press, & Smith, 2012) systematically varied task conditions in both the intrinsic and extrinsic spaces relative to each other, and suggest that the range of effects can best be accounted for by a model positing “gain field” representations spanning intrinsic and extrinsic space.

Details aside, the way to best summarize these results is as follows. In general, what determines generalization is *shared structure in the visuomotor space*. On a Bayesian perspective, this is best construed as adjusting one's forward model to account for the patterns of variation in the training set. But these representations are not detached from the visuomotor structure, as amodal representations would be. As Braun et al. (2009) note, the results suggest “that the feedback control process is not generic but is highly dependent on the structure that subjects had experienced within a single trial, which argues against a nonspecific increase in feedback adaptability” (354). That is, while the represented structures range over particular instances, that range is still represented in the same kind of space in which particular trials are performed.

4.3. Context.

The foregoing suggests that abstraction occurs within representations of visuomotor space. But one might reasonably ask how different actions are represented within the space. Since different actions require different generalizations, how are these “partitioned” (Howard, Ingram, & Wolpert, 2010) within a space that, presumably, is largely shared between different actions? One key principle that helps to achieve this partitioning is *context*. By analyzing different forms of context, investigators have explored the limits of generalization, and studied patterns of generalization and interference.

Howard et al. (2010) had subjects perform bimanual tasks, in which each hand reached to a target. In one context, both hands had to move in the same direction, while in the second context they had to move in opposite directions. Within each of these contexts, perturbations were introduced mid-trial. Perturbations could be in either of two opposite

directions. The study showed that when the two perturbations occurred in the same context, they interfered with each other, since they required precisely opposite adjustments. However, when each direction of perturbation was specific to a context, this interference drastically decreased. What this shows is that context – in this case the context of distinct bimanual actions – can partition learning and generalization within the sensorimotor space so that interference between tasks is reduced.

Sarwary et al. (Sarwary, Stegeman, Selen, & Medendorp, 2015) investigated context in a different way. In their paradigm, subjects performed a two-stage movement. The first movement would be towards a target, which could be in one of two directions from the starting point. After they reached the first target, a second target would appear, towards which the subjects would have to perform a second movement. During the second movement, the manipulandum was subjected to a force perpendicular to the required direction of movement. Importantly, the force could be introduced in either a clockwise or counterclockwise direction, and in the experiment each direction was paired to a specific first movement. The investigators showed that subjects quickly learned the force fields that were dependent on each initial movement, and that they could generalize within that context. Moreover, context-dependent facilitation varied continuously as the initial movement did – so, if the initial movement direction was shifted so that the difference between the two contexts was lessened, there was a continuous drop-off in facilitation effects. So, generalization applies both within a context and in the cues that signal context.

While in these examples the context cues are motoric, it has also been shown that visual cues can disambiguate, at least partially, between motor contexts (Krouchev & Kalaska, 2003). Intriguingly, there is even suggestion that this context-dependent generalization can scale up to types of tasks being performed in types of experimental setups (e.g., a key press behavior being performed in front of a monitor; see Memelink & Hommel, 2005, 2006). These results combined are further evidence that forward models are represented in sensorimotor space, and that the motor system can make use of contextual cues directly to decide *which* forward model is appropriate for the task.

4.4. Summary.

The question we set out with at the beginning of this section was whether motor control was best construed as occurring within representations of visuomotor space. The foregoing results suggest a resounding 'yes'. As we saw, views of motor control based on generative models explain control as the result of a forward model interacting with sensory feedback. What the results from perturbation, generalization, and context manipulations suggest is that those forward models operate within a representation of sensorimotor space, combining intrinsic and extrinsic coordinate frames in order to anticipate and react to the subject's changing environment. Since precisely the effects that generative models of motor

control are cited to explain occur within that space, it is unsurprising that the models should also take that form, and this is uniformly how they are interpreted in the literature. In the next section, I will consider what this means for broader questions about cognitive architecture.

5. Cognitive Architecture Again.

I have suggested that a broadly Bayesian approach is compatible with multiple different posits about the format and structure of mental models/representations. This in turn leaves open a *pluralist* position, on which the mind – even if it is broadly Bayesian in nature – employs different kinds of models for different domains (cf. Quilty-Dunn, 2018). Here I explore the possibility that the interfaces between different sorts of representations are just the points where the transitivity of specific prediction and error processing is broken. Hence, there is no need to solve interface problems, and we can think about the respective roles of different kinds of representations, vis-à-vis each other, in other ways.

Traditional views of action control presume two functional posits that, according to OFC, are false. The first is that decision is proprietary to representations in lexical format. We have seen that, on OFC, the motor system is also involved in deliberation and choice – it represents goals and selects the best means of acting on them. Second, traditional views assume that abstract goal representation, at least at some level of abstraction, is done in propositional format. But the results from structural learning and context representation from within the Bayesian framework suggest that this too is false. Through dimensional learning and representation, the motor system can form lower-dimensional models that capture the variation across specific task instances, and can partition types of actions via context-cues.

To summarize, these results suggest that the motor system does not *require* specific input from propositional representation to enact motor commands, and conversely that the motor system itself has the resources needed to control complex action in its forward models. If this is the case, then no translation from one kind of content to another is required. There is also functional reason to suggest that, even if it is not *impossible* to represent sensorimotor structures with discursive representations, the system would not go about representing them in that way. Each specific discursive representation is fixed in the level of abstraction at which it represents the world, because its semantics are fixed arbitrarily and not structurally (for further discussion, see Kulvicki, 2015; ([Author's paper])). Hence, it would be computationally demanding to generate a specific representation to specify a motor action in detail, and representations of ranges of sensorimotor situations would require several distinct representations.

This allows for heading off a particular objection that could be made in response to my analysis of the experiments above. It might be argued, from the continuous hierarchy perspective, that all of the examples I cited just happen to be ones on which prediction error can be explained away by a structural model at a relatively low level. That is, when error is detected in these experiments (e.g., through a perturbation), the system can adjust because the error signal falls within their represented structure, but other sources of error might require higher-level processing.⁶ Each of the models I have discussed deal with things like the position of objects and the uncertainty that things like perturbations or variance across trials can introduce. But, as Rescorla noted, these assume a goal representation. Sources of error with regard to the goal might in turn require response from higher-level models.

An old philosophical example can help make the challenge clear. Suppose I am reaching for a glass, with the goal of taking a sip of gin. If I find out the glass in fact contains petrol, then this, one presumes, is not something that will be represented within the kinds of models I have discussed, since it is not the position or perceptual qualities of the goal that have changed (“been perturbed”), but instead the categorical identity of the target.

There are two responses that the defender of the discontinuity view should make here. First, one can admit that some kinds of error signal would require higher-level processing, but deny that higher-level processing must equate to non-sensorimotor processing. For instance, some views of planning and decision-making posit that these processes rely on “event files,” which are pairings of actions and perceptual outcomes ([Author’s paper]). Representation of error at this level, due to the realization that the grasping action would eventuate in petrol-taste rather than water-taste, would be higher-level without having to leave the sensorimotor system.⁷

Second, one can turn the challenge around, and ask how the proponent of a continuity between modal and amodal processing in these cases deals with the interface problem. Higher-level sensorimotor models have the nice property of representing variations in structure in the same behavioral space as the lower-level models that govern the action. This allows them to guide lower-level models. But, as the interface problem states, it is unclear how amodal representations can guide those processes. Given that (i) abstract sensorimotor models are involved in decision-making, and (ii) that they represent goals, then there is no argument to be made from the requirement of higher-level error correction to the notion that amodal representations must be involved.

⁶ Thanks to an anonymous reviewer for pushing me to handle this objection explicitly.

⁷ It has also recently been argued that *categorical* perception can operate in structured representations ([Author’s paper]). If this is true, then even changes in categorical identity within objects can be, in many cases at least, recognized without leaving models in sensorimotor format.

If I am on the right track here, then we can begin to think about alternative roles that amodal representations might play in action control. These roles, I suggest, do not require precise prediction from one type of representation to the other, and hence are natural points at which the transitivity of processing posited by the continuous hierarchy view might be broken. I discuss two such roles: *target biasing* and *context cues*.

Take target biasing first. Results from the priming literature suggest that lexical representations have an incredible power to guide attention to objects, but this attention is of a rather diffuse sort – priming occurs towards broad dimensions rather than to specific outcomes (Memelink & Hommel, 2013; Weidner & Muller, 2013; Zehetleitner, Rangelov, & Müller, 2012). On the current perspective, guiding attention to targets can make it more likely that those outcomes will be acted towards by the motor system, without instructing the motor system precisely what it should do to act towards those objects. As we have seen, targets and trajectories must themselves be represented in the motor system. Cuing or biasing the motor system towards an outcome it already represents is distinct from determining the details of those representations.

A similar possibility exists with regards to cuing *contexts*. We have seen that contexts can be represented within the motor system and that generalization results occur within them. But a lexical cue can be one way in which a particular context model is activated. It is well established, for instance, that instructing subjects to prepare a grasping action primes them to recognize certain kinds of action targets (i.e., ones discriminated by size and orientation). One can imagine myriad other examples – for instance one’s tennis coach telling them to hit a forehand – which would simply serve to activate a model that the motor system has already learned or is in the process of learning. This would not require the lexical representation to stand in any specific semantic relationships to any of the details of the motor representation, but simply to index a task-relevant forward model within which the relevant motor control is enacted. Again, while compatible with a Bayesian approach, nothing about this entails that the discursive representation be able to predict in detail anything about what will go on in sensorimotor space. Hence, it is one reasonable juncture at which to posit a break in transitivity, and therefore a break in the continuous hierarchy.

6. Conclusion

One constantly touted advantage of a Bayesian/predictive processing approach to the mind is its unifying power. Those who like unification are impressed with the possibility that, at some level, all parts of the mind are doing the *same* thing, namely Bayesian updating. I have suggested that, even if the entire brain is Bayesian, this does not provide any motivation for positing a continuous cognitive architecture. Indeed, even a Bayesian brain may have different ways of representing the world, which it employs in particular functional domains. If so, then this changes the ways in which we theorize about the

interactions between distinct parts of the mind. Understanding cognitive architecture requires understanding its disunified aspects, as well as what all parts might share in common.

REFERENCES

- Anderson, M. L. (2017). Of Bayes and bullets: An embodied, situated, targeting-based account of predictive processing.
- Braun, D. A., Aertsen, A., Wolpert, D. M., & Mehring, C. (2009). Motor task variation induces structural learning. *Curr Biol*, *19*(4), 352-357. doi:10.1016/j.cub.2009.01.036
- Braun, D. A., Mehring, C., & Wolpert, D. M. (2010). Structure learning in action. *Behav Brain Res*, *206*(2), 157-165. doi:10.1016/j.bbr.2009.08.031
- Braun, D. A., Waldert, S., Aertsen, A., Wolpert, D. M., & Mehring, C. (2010). Structure learning in a sensorimotor association task. *PloS one*, *5*(1), e8973.
- Brayanov, J. B., Press, D. Z., & Smith, M. A. (2012). Motor memory is encoded as a gain-field combination of intrinsic and extrinsic action representations. *J Neurosci*, *32*(43), 14951-14965. doi:10.1523/JNEUROSCI.1928-12.2012
- Bruton, M., & O'Dwyer, N. (2018). Synergies in coordination: a comprehensive overview of neural, computational, and behavioral approaches. *Journal of neurophysiology*, *120*(6), 2761-2774.
- Burnston, D. C. (2017). Interface problems in the explanation of action. *Philosophical Explorations*, *20*(2), 242-258.
- Burnston, D. C. (2020). Anti-Intellectualism for the Learning and Employment of Skill. *Review of Philosophy and Psychology*. doi:10.1007/s13164-020-00506-5
- Burnston, D. C. (2020). Fodor on imagistic mental representations. *Rivista internazionale di Filosofia e Psicologia*, *11*(1), 71-94.
- Butterfill, S. A., & Sinigaglia, C. (2014). Intention and motor representation in purposive action. *Philosophy and Phenomenological Research*, *88*(1), 119-145.
- Cao, R. (2020). New Labels for Old Ideas: Predictive Processing and the Interpretation of Neural Signals. *Review of Philosophy and Psychology*, *11*(3), 517-546. doi:10.1007/s13164-020-00481-x
- Carroll, T. J., Poh, E., & de Rugy, A. (2014). New visuomotor maps are immediately available to the opposite limb. *J Neurophysiol*, *111*(11), 2232-2243. doi:10.1152/jn.00042.2014
- Cisek, P., & Kalaska, J. F. (2005). Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. *Neuron*, *45*(5), 801-814.
- Cisek, P., & Kalaska, J. F. (2010). Neural mechanisms for interacting with a world full of action choices. *Annu Rev Neurosci*, *33*, 269-298. doi:10.1146/annurev.neuro.051508.135409
- Colombo, M., Elkin, L., & Hartmann, S. (2018). Being Realist about Bayes, and the Predictive Processing Theory of Mind.
- Colombo, M., Elkin, L., & Hartmann, S. (2020). Being realist about Bayes, and the predictive processing theory of mind. *The British Journal for the Philosophy of Science*.

- Colombo, M., & Hartmann, S. (2017). Bayesian cognitive science, unification, and explanation. *The British Journal for the Philosophy of Science*, 68(2), 451-484.
- Colombo, M., & Series, P. (2012). Bayes in the Brain--On Bayesian Modelling in Neuroscience. *British Journal for the Philosophy of Science*, 63(3), 697-723.
- d'Avella, A., Saltiel, P., & Bizzi, E. (2003). Combinations of muscle synergies in the construction of a natural motor behavior. *Nature neuroscience*, 6(3), 300-308.
- Davidson, D. (1980). *Essays on actions and events: Philosophical essays*. Oxford: Oxford University Press.
- Dewhurst, J. (2017). Folk psychology and the Bayesian brain.
- Diamond, J. S., Nashed, J. Y., Johansson, R. S., Wolpert, D. M., & Flanagan, J. R. (2015). Rapid visuomotor corrective responses during transport of hand-held objects incorporate novel object dynamics. *Journal of Neuroscience*, 35(29), 10572-10580.
- Drayson, Z. (2017). Modularity and the predictive mind.
- Ferretti, G., & Caiani, S. Z. (2018). Solving the interface problem without translation: The same format thesis. *Pacific philosophical quarterly*.
- Franklin, D. W., & Wolpert, D. M. (2008). Specificity of reflex adaptation for task-relevant variability. *Journal of Neuroscience*, 28(52), 14165-14175.
- Friston, K. (2011). Embodied inference: or "I think therefore I am, if I am what I think".
- Friston, K., Samothrakis, S., & Montague, R. (2012). Active inference and agency: optimal control without cost functions. *Biol Cybern*, 106(8-9), 523-541. doi:10.1007/s00422-012-0512-8
- Friston, K., Thornton, C., & Clark, A. (2012). Free-energy minimization and the dark-room problem. *Frontiers in psychology*, 3, 130.
- Gladziejewski, P., & Milkowski, M. (2017). Structural representations: causally relevant and different from detectors. *Biol Philos*, 32(3), 337-355. doi:10.1007/s10539-017-9562-6
- Goodman, N. D., Tenenbaum, J. B., & Gerstenberg, T. (2014). *Concepts in a probabilistic language of thought*.
- Haruno, M., Wolpert, D. M., & Kawato, M. (2003). *Hierarchical MOSAIC for movement generation*. Paper presented at the International congress series.
- Hayashi, T., Yokoi, A., Hirashima, M., & Nozaki, D. (2016). Visuomotor Map Determines How Visually Guided Reaching Movements are Corrected Within and Across Trials. *eNeuro*, 3(3). doi:10.1523/ENEURO.0032-16.2016
- Hewitson, C. L., Sowman, P. F., & Kaplan, D. M. (2018). Interlimb Generalization of Learned Bayesian Visuomotor Prior Occurs in Extrinsic Coordinates. *eNeuro*, 5(4). doi:10.1523/ENEURO.0183-18.2018
- Hipolito, I., & Kirchhoff, M. D. (2019). The Predictive Brain: A Modular View of Brain and Cognitive Function?
- Howard, I. S., Ingram, J. N., & Wolpert, D. M. (2010). Context-dependent partitioning of motor learning in bimanual movements. *Journal of neurophysiology*, 104(4), 2082-2091.
- Kaplan, D. M., & Hewitson, C. L. (2020). Modelling Bayesian Computation in the Brain: Unification, Explanation, and Constraints. In *Neural Mechanisms* (pp. 11-33): Springer.
- Kiefer, A., & Hohwy, J. (2017). Content and misrepresentation in hierarchical generative models. *Synthese*, 195(6), 2387-2415. doi:10.1007/s11229-017-1435-7

- Kiefer, A., & Hohwy, J. (2019). Representation in the prediction error minimization framework. *Routledge handbook to the philosophy of psychology*. Routledge.
- Kirchhoff, M. (2018). The body in action: predictive processing and the embodiment thesis. In *Oxford handbook of cognition: embodied, extended and enactive* (pp. 243-260): Oxford University Press.
- Kirchhoff, M. D., & Robertson, I. (2018). Enactivism and predictive processing: A non-representational view. *Philosophical Explorations*, 21(2), 264-281.
- Körding, K. P., & Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, 427(6971), 244.
- Krouchev, N. I., & Kalaska, J. F. (2003). Context-dependent anticipation of different task dynamics: rapid recall of appropriate motor skills using visual cues. *Journal of neurophysiology*, 89(2), 1165-1175.
- Kulvicki, J. (2015). Analog representation and the parts principle. *Review of Philosophy and Psychology*, 6(1), 165-180.
- Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people. *Behav Brain Sci*, 40, e253. doi:10.1017/S0140525X16001837
- Lupyan, G., & Clark, A. (2015). Words and the World Predictive Coding and the Language-Perception-Cognition Interface. *Current Directions in Psychological Science*, 24(4), 279-284.
- Memelink, J., & Hommel, B. (2005). Attention, instruction, and response representation. *European Journal of Cognitive Psychology*, 17(5), 674-685.
- Memelink, J., & Hommel, B. (2006). Tailoring perception and action to the task at hand. *European Journal of Cognitive Psychology*, 18(4), 579-592.
- Memelink, J., & Hommel, B. (2013). Intentional weighting: a basic principle in cognitive control. *Psychological research*, 77(3), 249-259.
- Merel, J., Botvinick, M., & Wayne, G. (2019). Hierarchical motor control in mammals and machines. *Nat Commun*, 10(1), 5489. doi:10.1038/s41467-019-13239-6
- Mylopoulos, M., & Pacherie, E. (2017). Intentions and motor representations: The interface challenge. *Review of Philosophy and Psychology*, 8(2), 317-336.
- Mylopoulos, M., & Pacherie, E. (2018). Intentions: The dynamic hierarchical model revisited. *Wiley Interdisciplinary Reviews: Cognitive Science*, e1481.
- Overduin, S. A., d'Avella, A., Roh, J., Carmena, J. M., & Bizzi, E. (2015). Representation of muscle synergies in the primate brain. *Journal of Neuroscience*, 35(37), 12615-12624.
- Pacherie, E. (2008). The phenomenology of action: A conceptual framework. *Cognition*, 107(1), 179-217.
- Pavese, C. (2017). A theory of practical meaning. *Philosophical Topics*, 45(2), 65-96.
- Pavese, C. (2018). On the psychological reality of practical representation. *Philosophical Psychology*.
- Rescorla, M. (2016). Bayesian sensorimotor psychology. *Mind & Language*, 31(1), 3-36.
- Sarwary, A. M., Stegeman, D. F., Selen, L. P., & Medendorp, W. P. (2015). Generalization and transfer of contextual cues in motor learning. *Journal of neurophysiology*, 114(3), 1565-1576. Retrieved from <https://journals.physiology.org/doi/pdf/10.1152/jn.00217.2015>
- Quilty-Dunn, J. (2018). Perceptual pluralism. *Noûs*.

- Selen, L. P., Shadlen, M. N., & Wolpert, D. M. (2012). Deliberation in the motor system: reflex gains track evolving evidence leading to a decision. *J Neurosci*, *32*(7), 2276-2286. doi:10.1523/JNEUROSCI.5273-11.2012
- Sprevak, M. (2019). Two Kinds of Information Processing in Cognition. *Review of Philosophy and Psychology*, *11*(3), 591-611. doi:10.1007/s13164-019-00438-9
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science*, *331*(6022), 1279-1285.
- Todorov, E. (2006). Optimal control theory. *Bayesian brain: probabilistic approaches to neural coding*, 269-298.
- Todorov, E. (2009). Parallels between sensory and motor information processing. *The cognitive neurosciences*, 613-624.
- Todorov, E., & Ghahramani, Z. (2004, 1-5 Sept. 2004). *Analysis of the synergies underlying complex hand manipulation*. Paper presented at the The 26th Annual International Conference of the IEEE Engineering in Medicine and Biology Society.
- Vance, J., & Stokes, D. (2017). Noise, uncertainty, and interest: Predictive coding and cognitive penetration. *Consciousness and cognition*, *47*, 86-98.
- Vetter, P., & Newen, A. (2014). Varieties of cognitive penetration in visual perception. *Consciousness and cognition*, *27*, 62-75.
- Weidner, R., & Müller, H. J. (2013). Dimensional weighting in cross-dimensional singleton conjunction search. *J Vis*, *13*(3). doi:10.1167/13.3.25
- Williams, D. (2018a). Predictive coding and thought. *Synthese*, 1-27.
- Williams, D. (2018b). Predictive processing and the representation wars. *Minds and Machines*, *28*(1), 141-172.
- Williams, D. (2019). Hierarchical minds and the perception/cognition distinction. *Inquiry*, 1-23.
- Zednik, C., & Jäkel, F. (2016). Bayesian reverse-engineering considered as a research strategy for cognitive science. *Synthese*, *193*(12), 3951-3985. doi:10.1007/s11229-016-1180-3
- Zehetleitner, M., Rangelov, D., & Müller, H. J. (2012). Partial repetition costs persist in nonsearch compound tasks: Evidence for multiple-weighting-systems hypothesis. *Attention, Perception, & Psychophysics*, *74*(5), 879-890.
- Zhang, H., Daw, N. D., & Maloney, L. T. (2015). Human representation of visuo-motor uncertainty as mixtures of orthogonal basis distributions. *Nat Neurosci*, *18*(8), 1152-1158. doi:10.1038/nn.4055