
Of Bayes and Bullets:

An Embodied, Situated, Targeting-Based Account of Predictive Processing

Michael L. Anderson

Here I argue that Jakob Hohwy’s (Hohwy 2013) cognitivist interpretation of predictive processing (a) does not necessarily follow from the evidence for the importance of Bayesian processing in the brain; (b) is rooted in a misunderstanding of our epistemic position in the world; and (c) is undesirable in that it leads to epistemic internalism or idealism. My claim is that the internalist/idealist conclusions do not follow from predictive processing itself, but instead from the model of perception Hohwy’s adopts, and that there are alternate models of perception that do not lend themselves to idealist conclusions. The position I advocate is similar to Andy Clark’s embodied/embedded interpretation of Bayesian processing (Clark 2015); however, I argue that Clark’s position, as currently stated, also potentially leads to idealist conclusions. I offer a specific emendation to Clark’s view that I believe avoids this pitfall.

Keywords

Bayesian brain | Ecological psychology | Embodied cognition

1 Mechanical Applications of Bayes’ Theorem

As it is generally presented, Bayes’ theorem is a way of updating confidence in a given belief or hypothesis given accumulating evidence. In its modern mathematical form, it looks like this:

$$P(B|E) = [P(E|B)P_{\text{prior}}(B)] / [\sum P(E|B')P_{\text{prior}}(B')]$$

The probability P of belief B given evidence E is equal to the probability of E given B , times the previous estimate of the probability of B , divided by the sum of the probabilities of the evidence given all possible B s. Note that this formulation is fundamentally *representational*: it posits contentful beliefs about things in the world that are deemed more or less likely to be true in light of ongoing observations (Wiese and Metzinger 2017).

It is probably the case that this is the right interpretation to give those symbols when the equation is being used by human beings to determine things like the probability of a terrorist attack given intercepted phone calls, or a medical diagnosis in light of test results. But I think there is an important question to be asked about what interpretation to provide in cases where this mathematical relationship is being implemented by impersonal (or sub-personal) *mechanisms*. One of the earliest practical successes of Bayes’ Theorem was in artillery targeting, and a search of the military research literature suggests it remains in use today. The point for now is simple: automated targeting computers don’t believe anything, at least not yet, and still Bayes’ Theorem is effective at setting the parameters of the mechanism—getting the machine into the right configuration—to hit its target.¹

To signal my intentions at the outset, I believe that words like “targeting”, “(re)configuring” and “guiding” should serve as replacements for the prevailing metaphors of “belief”, “inference” and “hypothesis” at play in the majority of the predictive processing literature. Yes, we scientists can use Bayes’

¹ The details of the targeting mechanism are, it seems, classified. But roughly speaking the procedure involves feeding observations of munition landing positions (degree and direction of deviation from the target, along with several other variables) into a fire direction computer with a database of targeting settings/adjustments that were pre-computed using Bayesian methods. The output of the computer is relayed to the gunners who change settings including barrel angle and orientation accordingly. Note this *could* be done in a fully automated way, but for cultural reasons, the U.S. military prefers “human-in-the-loop” systems. See U.S. Army field manual 3-09.22 for the general outlines. Theoretical background can be found in Kolmogorov 1942, and Kolmogorov and Hewitt 1948. An interesting blog post on using Bayes to solve the related problem of artillery fire safety can be found here: <https://linesoftangency.wordpress.com/2012/01/11/check-one-two/>

rule to update our beliefs and hypotheses; it can also be used to appropriately change the parameters in a control system. I think it is an open question what Bayesian processing might be up to in our brains. It could be that it is building sub-personal representations including hypotheses, models and beliefs about a mind-independent world, and is thus an important part of the process of *perceptual reconstruction* of the shape of that world. But a more minimalist, mechanical interpretation of Bayesian updating is available to us, and we should take seriously the possibility that it offers a truer account of the Bayesian brain.

In what follows, I will argue that Hohwy's (Hohwy 2013) cognitivist interpretation of predictive processing, although consistent with the views of its most prominent proponents, (a) does not necessarily follow from the evidence for Bayesian processing; (b) is rooted in a misunderstanding of our epistemic position in the world; and (c) is undesirable in that it leads to epistemic internalism or idealism.

I will then turn to Clark's (Clark 2016) reading of the same literature, which purports to avoid the idealist pitfalls of Hohwy's account by situating predictive processing in the action-oriented framework offered by the embodied cognition movement (Anderson 2003; Anderson 2014; Clark 1997). Although I agree entirely with Clark's final analysis, I will suggest that in trying to combine aspects of both the cognitivist and the embodied accounts of predictive processing, Clark has generated a dilemma of perceptual content that ultimately lands him back in Hohwy's foundering epistemic boat.

I send Susan Oyama (Oyama 2000/1985) to the rescue. I argue that—at least at critical moments in his exposition—Clark (and Hohwy, too, although I won't make this case in detail) appears to be in the grip of an untenable, ultimately dualist account of the nature and origins of information, and this underlies his inability to escape epistemic internalism. I suggest that a more thoroughgoing dynamic systems account than the one Clark defends, along with a monist account of information according to which content emerges only in organism-environment interactions, offers both a way out of the internalist hole, and also a truer picture of the role of Bayesian updating in our cognitive economy. According to this picture, perceptual content is action-oriented, rooted in the fundamental intentionality of action² (Anderson and Chemero 2009; Anderson and Rosenberg 2008; Clark 2015), and Bayesian updating helps set the parameters that allow our dynamic, transiently assembled local neural sub-systems (TALoNS; Anderson 2014) to appropriately target adaptive behavior.

2 The “Problem” of Perception

It is commonly thought that the senses are inadequate for perception. It sounds odd, I hope, to put it that way, but this is indeed the majority position. For those who hold it, there's no point puzzling about how some three billion years of evolution could have failed to furnish us with better systems, because it is not just true but *obvious* that such is the case. After all, the world is 3-dimensional, and the retinal image not just flat, but inverted. The world varies along uncountable dimensions, and our senses respond to only a small fraction of them. On this view, the information delivered by our senses is at best an *impoverished reflection* of the world out there, and reconstructing that world from such materials takes both intelligence and wisdom. It takes intelligence, for instance, in the form of clever algorithms for extracting shape and texture from luminance cues, and wisdom in the form of knowledge that illumination generally comes from above. The brain that manages to solve the problem of perception does so just in so far as it is able to both *manipulate* and *add to* the information delivered by the senses.

2 Classically, intentional states are directed at things in virtue of a subject's perception of those things, and the intentional relation is grounded in the causal relation that gave rise to the perception (e.g. Devitt 1981; Grice and White 1961; Kripke 1980). On such views, the intentionality of action is derivative from the intentionality of perception and/or the intentional states guiding the action. In my view, this gets things exactly backward. Part of what it is to be an action is to be directed toward objects and ends; this intentionality is not derived from intentional mental states, but is inherent to behavior itself. Intentional mental states are directed at things *because* the actions they guide are, and in this sense the intentionality of mental states is derivative from the basic intentionality of action. Defending this view would of course require its own paper.

On the specifically Bayesian version of how the brain solves this problem of reconstruction, the trick is managed by using the deliverances of the sense organs as the basis for *inferences* about the ultimate causes of those stimulations. Beliefs about the world (explicit as well as tacit) both result from and constrain these inferences. As Hohwy expresses the point:

The problem of perception is a *problem* because it is not easy to reason from only the known effects back to their hidden causes. This is because the same cause can give rise to very different effects on our sense organs. [...] Likewise, different causes can give rise to very similar effects. (Hohwy 2013, p. 13)

This is a classic poverty-of-the-stimulus argument. Because the senses are incapable of specifying a unique situation, the reverse-inference problem posed by perception is too unconstrained to solve with sensory resources alone. Hohwy again: “[i]f the only constraint on the brain’s causal inference is the immediate sensory input, then, from the point of view of the brain, any causal inference is as good as any other.” (Hohwy 2013, p. 14) But of course, immediate sensory input is *not* the only constraint; there are, in addition, general beliefs about the world, specific hypotheses about the current state of the world, and *ongoing* sensory input. All this together, combined in the right way, is what on this view is needed to solve the problem of perception.

One other feature of the problem of perception, an unnecessary but oft included part of the story, takes the form of an additional metaphysical posit couched as a simple corollary. As Helmholtz put it: “[w]e always in fact only have direct access to the events at the nerves, that is, we sense effects, never the external objects.” (1867, p. 430; quoted in Hohwy 2013, p. 17) What this means is that in solving the problem of perception, in using the senses to answer the question of what’s out there, we are not permitted to “help ourselves to the answer by going beyond the perspective of the skull-bound brain.” (Hohwy 2013, p. 15) Thus, one of the things we need inference for is to “overcome the brain’s encapsulation in the skull” (Hohwy 2013, p. 41) and give us knowledge of what we cannot directly access.

One immediate point, worth making for our purposes here, is that predictive processing is only one proposed solution to what is supposed to be a very general problem. Very few of the apparent consequences for perception that follow from Hohwy’s treatment—that it is knowledge-rich, representational, and indirect, for instance—result from features specific to the predictive processing solution. Rather, these are consequences that would follow from *any* adequate solution to the problem as posed.

What is perhaps even more striking are the metaphors of containment that abound—that are apparently irresistible—when treating perception in this manner. We are “trapped inside our skulls” (Hohwy 2013, p. 224), “skull-bound” (pp. 15; 75), screened off from the world by a “veil of sensory input” (p. 50). Given this predicament, it’s the best we can do to infer the “hidden” (pp. 50; 81) causes; the only alternative would be to “impossibly jump outside the skull and compare percepts to the causal states of affairs in the real world.” (p. 50)

It’s worth asking exactly *who* is thus trapped? Who is the “we” inside the skull with “direct access [only] to the events at the nerves” and longing to “jump outside”? The answer has to be the knowing subject, the *person*, and although Hohwy is at pains to deny any hint of homuncularism, it’s hard not to notice the little man peeking out from behind that veil. It is no defense to point out—however correctly—that these are just metaphors. For as Lakoff and Johnson have been telling us for some time (e.g. Lakoff and Johnson 1980; Lakoff and Johnson 1999) metaphors *matter*. They change how we think, what we attend to, what questions we ask, and what answers we will accept.

The metaphors are certainly doing their work here. Hohwy’s initial statement of the problem of perception asks us to imagine being alone in a dark, sealed house and trying to determine what is making that tapping noise outside. It’s an effective image, in that it’s hard to conceive of resources *other* than reason that could be brought to bear in such restricted circumstances. It’s also wildly misleading. Sensory systems are nothing like that, passive, isolated, and unimodal (not to mention objective and reconstructive (Akins 1996; Anderson 2014)). Hohwy later admits the inadequacy of the metaphor

(p. 77) but by then the damage has been done. Although the house gets windows and legs, the walls remain intact, the doors are locked, and the mind (in this metaphor identical to the person) is trapped inside.

I confess, I don't recognize myself in this image. For a start, I'm pretty sure *I* encompass my skull; it is inside me, and not the other way around. It's a minor point, perhaps, but telling, for it forces us to question this notion of "access". Given that I am in the world (and not in my skull) in what sense am I cut off from that world? Do I not have access to air? To food? To sunlight and the solid surfaces that support me? If I have physical access to these things, what prevents having perceptual access?³ What is meant by the claim we have direct perceptual access only to events at the nerves? It's certainly the case that nerves lie along some important causal paths that run between organisms and the world. But what licenses the claim that our perceptual access to the world ends where the nerves do?

An obvious candidate answer would be that I have access only to the last link in the causal chain; the links prior are increasingly distal. But I do not believe that identifying our access with the cause most proximal to the brain can be made to work, here, because I don't see a way to avoid the path that leads to our access being restricted to the chemicals at the nearest synapse, or the ions at the last gate. There is always a cause even "closer" to the brain than the world next to the retina or fingertip. Moreover, causes aren't really linear chains, but complex webs of conditions. The causes of the shattering glass include the tension in my muscles, the energy they put into the descending hammer, the rigidity of the hammer's head, the brittleness of the glass, the firmness of the supporting surface, and more (see, e.g. Mackie 1965). In light of such complexity, it seems to me that identifying "the" causal link we access in perception is likely a fool's errand to start.

We are dangerously close to confusing epistemic and causal mediators here. Hohwy doesn't make this mistake, but he doesn't avoid it, either, for he slides between the *causal* claim that we have access only to events at the nerves (e.g. pp. 17; 50) and the *epistemic* claim that we have access only to sense data (e.g. p. 75). This slide introduces psychological and informational content where before there was only cause. It's not clear what warrants this, and, if it is warranted in general, what justifies the restriction of epistemic access to "sense data". If we have access to content and information, why conceptualize that access in so narrow a way? And if that content is *about the world* it is not clear what blocks the conclusion that we thereby access the world. On the other hand, if we have access only to causes, we are still owed a story about how those causes give rise to experience, and it's the details of this story that will tell us whether and how we have epistemic access to the world; epistemic encapsulation is not something to be declared by fiat. Hohwy knows (as Clark 2016, insists) that it is only in virtue of being embedded in the causal nexus that we have perceptual access to the world at all; far from isolating us from the world, our nervous system fully embeds us in it. What we stand behind in this picture is not a causal barrier, but an epistemic veil, and the image is arresting enough that we can fail to notice that its existence flows only from the metaphor of the house, and not from the facts of our situation.

As noted above, most of the conclusions I find objectionable here stem not from predictive processing per se, but from the account of what predictive processing must do to solve the problem of perception so formulated. In the next section, I'll briefly describe an alternative statement of the "problem" of perception that avoids these issues, and opens the path to a different treatment of predictive processing. But first we need to fill out Hohwy's account. For what is specific to Hohwy's interpretation of predictive processing is the story about exactly *how* beliefs, hypotheses, and sensory input get combined to enable perception. Arguably, that story exacerbates, or at least highlights, the apparent indirect nature of our perceptual—and hence epistemic—access to the world.

3 A critic might counter that the assumption that I am in the world, breathing air, is itself a mere seeming, compatible with being a brain-in-a-vat (which is indeed what Hohwy is arguing we are). Here I can only say that if I am not allowed the assumption that I am physically in the actual world, then it is the skeptic who is assuming what must be shown. More pointedly: it is the skeptic's way of setting up the problem of perception that opens up the gap between mind and world that (s)he fills with barriers and doubt. But this way is not the only way of understanding the job of perception. I sketch an alternative, below, from which these conclusions do not in fact follow (see also Anderson 2014; Anderson 2006).

The predictive processing story, neutrally specified, goes like this: our brains implement a massive, complex, hierarchical Bayesian network structured such that higher levels dynamically predict the states of lower levels, all the way down to the changing states of our sensory receptors and physical actuators. On this story, what flows “down” the hierarchy are predictions of state, and what flows “up” the hierarchy are the differences between the predicted and actual values. These error signals cause adjustments to the settings of the higher level(s) so as to bring the predictions into closer alignment with the actual states of the lower levels of the hierarchy.

It is universally acknowledged that this picture of the mechanisms of perception differs significantly from the classical model, which posits a feed-forward cascade of feature detection. There is, however, significant disagreement over what this difference in mechanism implies for the nature of perception itself (Butz 2017; Clark 2015; Clark 2016; Gallagher and Bower 2014; Hohwy 2013; Seth 2014). In Hohwy’s case, because he accepts the assumptions driving the classical account that the purpose of perception is to build objective models (representations) of the mind-independent world, his model of the cortical hierarchy is essentially identical to the classical model: perception results in the representation of properties of the environment at ever-higher levels of complexity and organization. As he puts the matter:

The brain responds to ... causal, hierarchical structure [in the world] in a very comprehensive manner: it recapitulates the interconnected hierarchy in a model maintained in the cortical hierarchy of the brain. (Hohwy 2013, p. 28)

I have no intention to mount a sustained argument against this position. I’ll simply make two points: First, there is nothing in the nature of the mechanism itself that warrants this interpretation. What a given layer of the cortical hierarchy is doing (and what kind of representations it trucks in) is an *empirical* question. It is not inferable from knowledge that the layer is performing Bayesian updating, any more than it would be inferable from knowledge that it is engaged in Hebbian plasticity. The brain is composed of functionally differentiated parts; no general mechanisms can tell us about these differences. Moreover, Hohwy’s interpretation is hardly a natural one if the basic idea is that the brain represents what it predicts; given that each layer is predicting the states of lower levels of the hierarchy, the hypothesis that sticks closest to the facts of the mechanism is that each layer represents the one below it (O’Regan and Degenaar 2014). More generally, I think there’s an interesting and unexamined commitment to two things, here, both of which are highly questionable. The first is structuralism about experience: the classic thesis that conscious experience is composed of simple *elements* (such as sense data), combined into different complex structures (Titchener 1929), and the second is an expectation that there must be an isomorphism between the structure of experience and the structure of the mechanism that gives rise to it. Clark (Clark 2016) is not entirely unaffected by these tacit commitments.

Second, once one combines a classical representational account of perception with a predictive processing mechanism, I believe that, although there may be defenses that could be mounted to resist the conclusion⁴, it is at least plausible that epistemic internalism naturally follows. Because the information from the world, flowing up the hierarchy, encodes only the mismatch between the prediction and reality, one’s sense of that reality can only be encoded in the prediction. As Hohwy succinctly states: “Perceptual content is determined by the hypothesis that best suppresses prediction error.” (Hohwy 2013, p. 117) On this model, perception is only indirectly caused by the world.

I most certainly do *not* want to recapitulate the realism vs. idealism and epistemic externalism vs. internalism debates that raged around the end of the 20th century (but see O’Donovan-Anderson 1997). I’ll simply assert without argument that the realists and externalists won, and that any episte-

⁴ I gestured at some possible options above.

mology that implies internalism and its attendant skepticism has made a mistake *somewhere*. I believe this conclusion is reliable enough that it can be used as a *reductio ad falsum*, and that is how I will employ it here: if internalism, then not (predictive coding & classical representationalism)⁵. Because we are taking predictive coding as true by hypothesis, it is classical representationalism that must go.

In this section, I have outlined Hohwy's understanding that the task of perception is to internally reconstruct the state of the world from impoverished sensory stimuli. For him, perception is a process of reverse inference, wherein the causes of sensory stimulation are inferred from their effects on sensory receptors. Because there is no unique solution to this problem—multiple different models could, in theory, account for the effects equally well—we are in the position of generating uncertain *hypotheses* about the true state of the world. According to Hohwy, predictive processing is the main neural-psychological mechanism we have for continually updating our uncertain models in light of ongoing stimulation. Our models/hypotheses make predictions about incoming states, and are updated in light of any mis-matches between the models and the states.

Hohwy argues that one important philosophical consequence of having such a perceptual system is that we have only indirect epistemic access to the actual world; he thus embraces a version of epistemic internalism that, under at least some circumstances, can lead to various forms of skepticism. Although I take issue with some of the steps in this argument, in the end I find the conclusion both plausible and unacceptable. To resist this conclusion, then, means denying one or more of its antecedents. For me, the obvious candidate is Hohwy's characterization of the problem of perception. I argue (1) that the internalist/idealist conclusions follow only from that model of perception, and not from predictive processing per se and (2) that there are alternate models of perception that do not lend themselves to idealist conclusions. It is to one such alternative model that I turn in the next section.

3 What Perception Is for⁶

One key aspect of the problem of perception as sketched above is that there is insufficient information delivered by the senses to specify the world. Another is the notion that the immediate function of perception is the veridical, objective representation of the external world. These two suppositions work together to support an *inferential, reconstructive* account of perception that centrally features the maintenance of world models. The core of Hohwy's account of predictive processing is an explanation of how Bayesian updating enables the construction and maintenance of such models. But what I want to suggest here is that both of these suppositions are questionable, if not demonstrably false, and therefore we don't need such an explanation. Rather, what needs explanation is how organisms appropriately attune to their environments to support adaptive behavior.

I'm not sure I know exactly what sense data are, but if a sense datum is meant to be something like the information about hue and intensity delivered by a single pixel of a digital camera, then I can say with confidence that neither the retina nor indeed any sensory system delivers sense data. Visual input to the brain is not like the snapshot of a camera, but involves multiple, mutually informing structured flows coming from the eyes (and not just the retina!), ears, head, limbs, and the rest of the active, sensing body. The raw materials of perception are not the momentary impacts of light on the retina, or chemicals on the olfactory receptors, but rather the relationships between *changes* across multiple modalities as one's position and posture change. As Gibson ([Gibson 1966](#)) writes:

⁵ In fact, this formulation is too simple, for clearly there are forms of representationalism that do not imply idealism/epistemic internalism. So in fact there are at least two choices for resisting Hohwy's conclusions: (1) reanalyze the origins of perceptual content in predictive processing while denying that epistemic access is restricted to sense-data, but otherwise accepting a representationalist framework; or (2) rejecting the representationalist gloss on predictive processing. I think the former is a live possibility that deserves sustained exploration, but I will follow out the latter here, in part because I think Andy Clark can be understood as offering a proposal of the former sort, and, as I will argue below, that proposal suffers from its own epistemic issues.

⁶ Portions of this section are adapted from chapter 5 of [Anderson 2014](#).

The active observer gets invariant perceptions despite varying sensations. He perceives a constant object by vision despite changing sensations of light; he perceives objects by feel despite changing sensations of pressure; he perceives the same source of sound despite changing sensations of loudness in his ears. The hypothesis is that constant perception depends on the ability of the individual to detect the invariants, and that he ordinarily pays no attention whatever to the flux of changing sensations. (Gibson 1966, p. 3)

When we attend to the fact that perception is an activity, that part of seeing, and smelling, and feeling is *moving*, we see that the actual deliverances of perception are extremely rich, multimodal, and perfectly capable of revealing the higher-order invariants in our environment and uniquely specifying the shape of the world. More to the point, perception doesn't start with sensory stimulation, for each and every "stimulation" was itself preceded (and is generally accompanied) by an action—action and perception are constant, ongoing and intertwined. As Gibson writes, "The active senses cannot simply be the initiators of signals in nerve fibers or messages to the brain; instead they are analogous to tentacles or feelers" (Gibson 1966, p. 5).

On this alternate view, then, the problem of perception is *not* how organisms get from stimulus to model, for perception and action are deeply linked. It is because one's view changes as eyes, head, and body move around the world that it is possible to know the world. This is why Hohwy's darkened house metaphor rigs the game: the perceptual system is an exploratory and not an inferential system. In active perception, the presumed poverty of the stimulus disappears. There is no infinity of possible worlds to sort through, each equally consistent with incoming information, for the deliverances of perception are not limited to the momentary 2-dimensional image on the retina (or the surface of the skin), but consist rather in the much richer set of distinctive transformations of that input given changes in posture and position. These transformations are sufficient to reveal the shape of the world.

Indeed, as Warren (Warren 2005) points out, the ability to pick up on such environmental invariances is likely a condition of perceptual systems evolving at all:

Perceptual systems become attuned to informational regularities in the same manner that other systems adapt to other sorts of environmental regularities (such as a food source): possessing the relevant bit of physiological plumbing (whether an enzyme or a neural circuit) to exploit a regularity confers a selective advantage upon the organism. Since the water beetle larva's prey floats on the surface of the pond and illumination regularly comes from above, possession of an eye spot and a phototropic circuit can enhance survival and reproductive success. But if illumination were ambiguous and prior knowledge were required to infer the direction of the prey, it is not clear how such a visual mechanism would get off the ground. Natural selection converges on specific information that supports efficacious action.

What the [traditional] view treats as assumptions imputed to the perceiver can thus be understood as ecological constraints under which the perceptual system evolved. The perceptual system need not internally represent an assumption that natural surfaces are regularly textured, that terrestrial objects obey the law of gravitation, or that light comes from above. Rather, these are facts of nature that are responsible for the informational regularities to which perceptual systems adapt, such as texture gradients, declination angles, and illumination gradients. They need not be internally represented as assumptions because the perceptual system need not perform the inverse inferences that require them as premises. The perceptual system simply becomes attuned to information that, within its niche, reliably specifies the environmental situation and enables the organism to act effectively. (Warren 2005, pp. 357-8)

This brings us to a second supposition that underlies the traditional approach. Insofar as it is the fundamental job of the perceptual system to build and maintain a model of the world, then it becomes

natural to think the fundamental content of perception is written in a perceiver-independent physics-influenced vocabulary of edges, colors, velocities and weights (Gibson 1979). But if the purpose of perception is to guide action, then the more parsimonious hypothesis is that the organism will be primarily sensitive not to these sorts of properties but rather to action-relevant relationships between organism and environment (Anderson and Rosenberg 2008; Anderson and Chemero 2009).

The frog's visual system, for example, is tuned to particular patterns of motion that, in the restricted context of its niche, specify small edible prey and large looming threats. The fish's lateral line organ is tuned to pressure waves that specify obstacles, the movements of predators and prey, and the positions of neighbors in the school. Even the narwhal's tusk turns out to be a sense organ tuned to salinity differentials that specify the freezing of the water's surface overhead. The narwhal is thereby in perceptual contact with a property of its niche—the penetrability of the surface—that is critical to its survival. (Warren 2005, pp. 340-1)

It is the overall job of perceptual systems to keep organisms in contact with the values of relevant organism-environment relationships (the closeness of the obstacle; the penetrability of the surface). Put differently, the world properties it is important to pick out for the purpose of *reconstruction* are not the same as those that best support *interaction*, and psychology has tended to (mistakenly) focus on the former class of properties to the exclusion of the latter.

This, of course, is the thought behind the Gibsonian affordance-based theories of perception that have been widely influential in embodied cognition (Gibson 1979; Orlandi 2014). Affordances are relationships between things in the world and an organism's abilities (Chemero 2009): for the average human the chair (but not the twig) affords sitting, the path (but not the wall) affords walking, and the rock (but not a tiny urticating bristle) affords throwing, but things are different for (and look different to) the bird and the spider. According to the view I am advocating here, perception is primarily perception of such affordances; the world is seen as a changing set of opportunities for action and interaction. Perception is not for building *models* of the world; it is for building *control systems* for the organism.

Cisek (Cisek 1999) has developed this line of thought in a deeply interesting way:

As evolution produced increasingly more complex organisms, the mechanisms of control developed more sophisticated and more convoluted solutions to their respective tasks. Mechanisms controlling internal variables such as body temperature or osmolarity evolved by exploiting consistent properties of chemistry, physics, fluid dynamics, etc. Today we call these “physiology”. Mechanisms whose control extends out through the environment had to exploit consistent properties of the environment. These properties include statistics of nutrient distributions, Euclidean geometry, Newtonian mechanics, etc. Today we call such mechanisms “behavior”. In both cases the functional architecture takes the form of a negative feedback loop, central to which is the measurement of some vital variable. Fluctuations in the measured value of the variable outside some “desired range” initiate mechanisms whose purpose is to bring the variable back into the desired range... The alternative “control metaphor” being developed here may now be stated explicitly: *the function of the brain is to exert control over the organism's state within its environment.* (Cisek 1999, pp. 8-9. emphasis in original)

The beauty of Cisek's analysis here is that he shows how the idea of (negative) feedback loops, which are known to be of crucial importance to the regulatory systems of living things at multiple spatial scales and levels of analysis, can be generalized to cover the case of overt behavior. If that's a valid reconceptualization, and I think it is very, very promising (Anderson 2014), one effect should be to shift our focus from how brain mechanisms like Bayesian predictive coding implement and maintain models of the world, to how such mechanisms enable the feedback loops that maintain attunement to the environment and support adaptive behavior.

4 Clark and Coding

It would appear that Clark (Clark 2016) supports this shift of focus. There is evidence for that thought:

By the end of our story, the predictive brain will stand revealed not as an isolated inner ‘inference engine’ but an action-oriented engagement machine—an enabling ... node in a pattern of dense reciprocal exchange binding brain, body and world. (Clark 2016, p. xvi)

Consistent with the research program I advocated for at length in Anderson (Anderson 2014), Clark suggests that we should understand the search of possibility space that Bayesian predictive processing enables is not the search for the best hypothesis or world model, but rather the search for the best sensorimotor machine:

Integration (of the rather profound kind exhibited by the neural economy) means that those functionally differentiated areas [of the brain] interact dynamically in ways that allow transient task-specific processing regimes (including transient coalitions of neural resources) to emerge as contextual effects [that] repeatedly reconfigure the flow of information and influence. [...] It is the guidance of world-engaging action, and not the production of ‘accurate’ internal representations, that is the real purpose of the prediction error minimizing routine. [...] [It] must find the set of neuronal states that best *accommodate* (as I will now put it) the current sensory barrage. (Clark 2016, pp. 142; 168; 192. Emphasis in original)

And yet, he also endorses positions that would seem perfectly at home in Hohwy’s book:

Perception is controlled hallucination. [...] active agents get to structure their sensory flow [...] [b]ut it remains correct to say that that all the system has direct access to is its own sensory states (patterns of stimulations across its sensory receptors. [...] The task is [...] to infer the nature of the signal source (the world) from just the varying input signal itself. [...] The ongoing process of perceiving, if such models are correct, is a matter of the brain using stored knowledge to predict, in a progressively more refined manner, the patterns of multilayer neuronal response elicited by the current sensory stimulation. This in turn underlines the surprising extent to which the structure of our expectations (both conscious and non-conscious) may be determining what we see, hear, and feel. (Clark 2016, pp. 14; 16; 27)

I find myself confronted by deep tensions in Clark’s account. On the one hand we have a thoroughgoing action-oriented, affordance-laden, sensorimotor coupling account of an agent’s ongoing engagement with the world. Bayesian predictive processing emerges here as a crucial mechanism for modulating the agent-environment coupling by dynamically adjusting the parameters of the neural mechanisms that support the agent’s capacity to target its behaviors. According to this story, “prediction-driven learning delivers a grip on affordances” (Clark 2016, p. 171) and thereby reveals—allows the perception of—a specifically human world (Clark 2016, p. xv). Here we are thoroughly, actively *in* the world, and it would appear to be impossible (or at least extremely unnatural) to conclude that we are epistemically isolated from it.

And yet we also have, on the other hand, the persistence of the model of perception not as attunement, but as inference from effects to causes, which leads to metaphors of cognitive confinement, of perception as hallucination (p. 14), as an *Ender’s Game*-style simulation in which our access to reality is mediated by a virtualization of it (p. 135). As nice as it would be for our expository task here if the inferential model of perception could be understood as one expository step on the way to understanding, to be discarded as enlightenment is attained, in fact the two models of perception remain in tension quite late in the book. Consider the following:

[P]rediction errors help select among (while simultaneously responding to) competing higher-level hypotheses, each of which implies a whole swathe of sensory and motor predictions. Such high-level hypotheses are intrinsically affordance-laden. (Clark 2016, p. 187)

Elsewhere on that same page, Clark suggests that neural mechanisms also select “salient representations that have affordance [i.e.] sensorimotor representations that predict both perceptual and behavioral consequences.” (p. 187) Note the slide from getting a “grip” on affordances (in the world) to *representing* them (in the head). Although Clark is at pains to deny that such representations act as epistemic mediators (p. 195), it is pretty clear we’ve taken at least a half step back inside the mind. Why? Because Clark cannot quite shake the influence of the inferential model of perception. Discussing the epistemic internalism adopted by many advocates of the predictive processing framework, he writes:

There is something right about all this and something (or so I shall argue) profoundly wrong. What is right is that accounts on offer depict perception as *in some sense* an inferential process: one that cannot help interpose *something* (the inference) between causes (such as sensory stimulations or distal objects) and effects (percepts, experiences). (Clark 2016, p. 170, emphasis in original)

I’m not quite sure how to analyze the tension in Clark’s account, but I suspect it is driven by his quasi-neo-Kantian account of perception, according to which experience is not the immediate apprehension of the material world, but sensation transformed by the subject. Hohwy unabashedly accepts that perceptual content is driven by top-down expectations, but Clark is trying to avoid the internalist consequences of that view by providing a rather greater epistemic role for the incoming sensory signal. Like the Kantian percept that draws concepts into operation, Clark’s sensory signals “select” the representations that determine the content of perceptual experience. Moreover, he suggests that increasing the gain on the sensory signal is epistemically equivalent to changing the amount or degree of content that sensory information drives:

This weighting determines the balance between top-down expectation and bottom-up sensory evidence. That same balance, if the class of models we have been pursuing is on track, determines what is perceived and how we act. (Clark 2016, p. 221)

Just as Kant attempted to avoid Humean skepticism by insisting on the necessity to experience of both percept *and* concept, so Clark tries to avoid Hohwian idealism by insisting on the importance to perception of both sensation and prediction. It is because of the balancing act between bottom-up and top-down signals that Clark suggests that we should say we have “not-indirect perception” (Clark 2016, p. 195). The double negative speaks volumes about Clark’s conceptual struggle, here. Now, neither Clark nor Hohwy actually offer a detailed theory of content (and, indeed, there has been little work along these lines, but see Gładziejewski 2016, and Wiese 2016), so my analysis is an admittedly speculative attempt to make sense of the tension in Clark’s account. But where Hohwy can probably avail himself of a fairly standard account of narrow content (e.g. conceptual role semantics like Loar 1988; Chalmers 2002; or a more radical internalism such as Segal 2000), it is very unclear what theoretical resources Clark has to work with. This is because his externalism would seem to require the sensory signal to carry content (perhaps to be analyzed along the lines of a causal theory of content, e.g. Dretske 1981), but he at the same time accepts Hohwy and Friston’s contention that top-down expectations carry (and at least partly determine) perceptual content. From this emerges the notion that these two sources of content must be “balanced” and together determine the content of experience.

I do not know whether or how this can be made to work. If we admit, as Clark appears to do, that the top-down prediction gives us only *indirect* knowledge of the world, then the direct access would have to be provided by the bottom up signal, perhaps by having content determined by the external cause of the signal. But by hypothesis the bottom-up sensory input is an *error* signal, specifying not

the world itself, but its deviation from predictions. Because the epistemic role played by the error signal is updating predictions/expectations, it appears that its epistemic effects can *only* ever be indirect (whatever its causal origins), leaving the top-down expectations to directly determine the content of perceptual experience (more on this in the next section). Although I am in complete agreement with Clark's *goals* in offering an externalist, embodied, action-oriented account of predictive coding, I must reluctantly conclude that he has not in fact fully succeeded in avoiding some undesirable epistemic consequences of the view. At the very least, building an externalist theory of content that is consistent with the whole of Clark's account faces significant challenges.

In the next section, I will suggest that the way forward involves denying that *either* the top-down or the bottom-up signal carry information, and adopting not a causal but a guidance theory of content, according to which content is determined by the intentional directedness of action, and not by the causal origins of perceptions.

5 Structure, Information, and Bayesian Mechanism.

There is perhaps no term in the cognitive sciences that is more abused than “information”. There is information in the world, in the brain, in the genes, in the sensory signal...information is ubiquitous, and so are scientific references to it. “Information” does an immense amount of work for us. However, if Oyama (Oyama 2000/1985) is right, and I think she is, this is a serious problem. I do not have space here to offer a complete account of her argument, but by way of illuminating my motivations for adopting her fairly radical solution of simply *denying* that there is information in any of these places, I believe she convincingly demonstrates that what she labels “preformationism”, the notion that information exists before its utilization or expression, lies behind some of our more persistent and pernicious dualisms, including mind-body, nature-nurture, person-situation, and nativism-empiricism, just to name a few. I would add internalism-externalism to that list, insofar as elements of that debate can seem to turn on determining just where content-specifying information comes from: the mind or the world?

Her argument is complex and subtle, but we can get most of the way to her conclusions by reflecting on the fact that a message has a meaning only on the assumption of a specific receiver or decoder. If different consumers extract different meanings from the same message, then the notion that it carries “information” is, if not false, then essentially useless for a theory of perceptual content. Insofar as a signal can carry the same e.g. Shannon-Weaver information relative to some specific situation, but trigger the instantiation of different semantic properties⁷ at different times or in different people, it makes little sense to think of those properties as being “in” the signal or information. If information is not inherently meaningful, supposing it specifies content is a mistake. It isn't even clear that one can substitute “structure” for “information” here, for although signals and messages are certainly and necessarily structured, the notion that it has “a” structure that thereby fixes content falls to the same considerations that deny it can be said to have “a” meaning: for different decoders, different aspects of its structure may be relevant.

What I want to argue here is that such considerations, along with my analysis of the purpose of perception offered above, should motivate us to give up on information-processing accounts of mind in favor of a developmental systems account. Central to such an account is the notion of *mutual constraint* between multiple interacting influences, at multiple spatial and temporal scales of analysis. Oyama develops the idea most fully in an account of gene-x-environment interactions in development, but I believe we can apply it to mind-world interactions as well. The root of Oyama's investigative problem is the question of whether organism-specifying information is in the genome, the environ-

⁷ Thanks to Thomas Metzinger for this way of putting the matter. Note that when Oyama denies there is information, she is really denying that there is content-specifying or outcome-specifying or otherwise inherently significant information. I do not believe she would deny (and I myself certainly do not) that any given signal can be said to have Shannon-Weaver information, but she would rightly point out that such information is specified *relative to a sender, receiver, and situation*.

ment, or both. Consider a gene that leads to an eye when expressed in its normal milieu, but that leads to a limb when expressed in a bodily location that normally produces neurons. What sense can be made of this? Surely the conclusion must be that neither the gene nor the environment specified what we thought when we observed normal development. Multiply this example by the myriad observations of such variability of phylogenetic outcome and it is hard to resist the conclusion that neither specifies any phylogenetic trait at all.

Similarly, one can ask where the content of experience comes from, the mind, the world, or both? Presumably, everyone will agree that the answer must be both, for we are not the passive recipients of the world's imprint, but active epistemic agents, shaping, categorizing and otherwise selecting incoming percepts. Yet, as in the case of gene-x-environment interactions, the 2-source solution is inherently unstable. Consider: if the error signal specifies a world-state, there isn't a need for the predictive model to specify that state (and we have epistemic support for a kind of passive externalism). If the error signal does *not* specify a world-state, then it must be the predictive model that does, and we oscillate back to internalism. But if content is somehow irreducibly determined by both, such that it is different from what either specifies on its own (or, to put it somewhat differently, if what each specifies or contributes to the whole depends on the state of the other) then it makes little sense to say that either intrinsically specifies anything at all.

On this sort of view, perceptual content is determined not by top-down expectations, nor by the incoming sensory stream, nor, Clark's solution, by both in varying degree. All these solutions have in common the preformationist fallacy that it is possible to specify the contribution of each interacting element in light of the information it brings to the interaction. This is what Oyama is at pains to deny. In systems marked by mutual causal constraint, what Oyama sometimes calls reciprocal selectivity, it is not generally possible to parse this out. Instead, we should say that perceptual experience is determined by the mutual constraint between the incoming sensory signal and ongoing neural and bodily processes, and no aspect of that content can be definitively attributed to either influence. As Oyama expresses the point:

A structured system selects its stimulus—indeed, defines it and sometimes produces it (the state of the system determines the kind and magnitude of stimulus that will be effective, and intrasystemic interactions may trigger further change)—and the stimulus selects the outcome (the system responds in one way rather than another, depending on the impinging influence). Nativists have generally focused on the former, while empiricists have stressed the latter. In doing so, they have perpetuated and further polarized the opposition between fated internal structure and fortuitous outside circumstance. The mutual selectivity of stimulus and system applies to causal systems of all sorts, and illustrates the impossibility of distinguishing definitively between internal and external control, the inherent and the imposed, selection and instruction. (Oyama 2000/1985 p. 68)

What an ongoing interaction *means* to an organism depends on the state of that organism, and is expressed in what the stimulus (and the organism) *does*. The sensory stimulation does not carry information independent of its causal effects, and these effects irreducibly depend on the state of the system within which the stimulation is occurring. From the standpoint of the project of this essay, what that means is that we should offer *only* a causal, but not an epistemic, information-processing, or representational gloss on the different roles of the top-down and bottom-up signal in predictive coding. Perceptual content is determined by the mutual constraint imposed between the interacting elements, which itself depends (and here I am in 100% agreement with Clark) on “transient task-specific processing regimes (including transient coalitions of neural resources) [that] emerge as contextual effects repeatedly reconfigure the flow [...] of influence.” (Clark 2016, p. 142) Put differently (and here I offer the seed of a new analysis of perception), *structure* (in the sensory inputs, inner states, and neural and bodily processes) may be intrinsic, but content is not. Perception is an event characterized by

the co-determination of an ongoing behavioral process by both inner and outer conditions. Structure becomes information via such events.

And how should we analyze the content of experience that dynamically emerges in these interactions? I believe that the guidance theory (Anderson and Chemero 2009; Anderson and Rosenberg 2008) is well suited to the job. According to the guidance theory, the intentionality of content rests on the fundamental intentionality of action. Put differently, content *is* what content *does*, and what it does is provide guidance for action. A full formalization of the theory is offered in (Anderson and Rosenberg 2008), but roughly speaking, a percept P (or representation R) is *of* an entity E just in case P (or R) is used to guide an agent's action with respect to E.⁸ On the view I am developing here, then, Bayesian updating should be understood as one crucially important neural process that gets the parameters of the mechanism properly set to guide behaviors to their targets.

6 Conclusion

In this essay I have attempted to trace the underlying causes of Hohwy's internalism to a faulty conception of the nature of our epistemic situation. I outline an alternative, embodied, situated and action-oriented perspective that should allow us to take predictive coding on board, while avoiding its internalist consequences. I then puzzle over the fact that, although Clark is fully on board with the action-oriented perspective I advocate, he hasn't quite entirely avoided internalism. I suggest it may be because he has replaced Hohwy's one-content-source account of perception with a two-content-source view, when he should have instead rethought the epistemic interpretation of predictive coding entirely. I then briefly offer a *merely causal* account of predictive coding, and gesture at a theory of perceptual content that fits nicely with the action-oriented perspective developed here.

⁸ The guidance theory is neutral and pluralist on the question of *how* guidance representations serve their representational function of guiding action. Guidance representations can be map-like, or emulators, or pictures, or something else entirely, and we expect different kinds of representations to be used in different circumstances and by different mental systems/subsystems.

References

- Akins, K. (1996). Of sensory systems and the aboutness of mental states. *The Journal of Philosophy*, 93, 337–372.
- Anderson, M. L. (2003). Embodied cognition: A field guide. *Artificial Intelligence*, 149 (1), 91–130.
- (2006). Cognitive science and epistemic openness. *Phenomenology and the Cognitive Sciences*, 5 (2), 125–154.
- (2014). *After phrenology: Neural reuse and the interactive brain*. Cambridge, MA: MIT Press.
- Anderson, M. L. & Chemero, A. (2009). Affordances and intentionality: Reply to Roberts. *Journal of Mind and Behavior*, 30 (4), 301.
- Anderson, M. L. & Rosenberg, G. (2008). Content and action: The guidance theory of representation. *Journal of Mind and Behavior*, 29 (1–2), 55–86.
- Burr, C. & Jones, M. (2016). The body as laboratory: Prediction-error minimization, embodiment, and representation. *Philosophical Psychology*, 29 (4), 586–600. <https://dx.doi.org/10.1080/09515089.2015.1135238>.
- Chalmers, D. (2002). The components of content. In D. Chalmers (Ed.) *The philosophy of mind: Classic and contemporary readings* (pp. 607–633). Oxford, Oxford University Press.
- Chemero, A. (2009). *Radical embodied cognitive science*. Cambridge, MA: MIT Press.
- Cisek, P. (1999). Beyond the computer metaphor: Behaviour as interaction. *Journal of Consciousness Studies*, 6 (11–12), 125–142.
- Clark, A. (1997). *Being there: Putting brain, body, and world together again*. Cambridge, MA: MIT Press.
- (2015). Predicting peace: The end of the representation wars. In T. K. Metzinger & J. M. Windt (Eds.) *Open MIND*. <https://dx.doi.org/10.15502/9783958570979>. <http://open-mind.net/papers/predicting-peace-the-end-of-the-representation-wars>.
- (2016). *Surfing uncertainty: Prediction, action, and the embodied mind*. New York: Oxford University Press.

- Devitt, M. (1981). *Designation*. New York: Columbia University Press.
- Dretske, F. (1981). *Knowledge and flow of information*. Cambridge, MA: MIT/Bradford Press.
- Gallagher, S. & Bower, M. (2014). Making enactivism even more embodied. *AVANT*, 5 (2), 232–247.
- Gibson, J. J. (1966). *The senses considered as perceptual systems*. Boston: Houghton-Mifflin.
- (1979). *The ecological approach to visual perception*. Hillsdale, NJ: Erlbaum.
- Grice, H. & White, A. (1961). Symposium: The causal theory of perception. *Proceedings of the Aristotelian Society, Supplementary Volumes*, 35, 121–168. <http://www.jstor.org/stable/4106682>.
- Gładziejewski, P. (2016). Predictive coding and representationalism. *Synthese*, 559–582. <https://dx.doi.org/10.1007/s11229-015-0762-9>.
- Hohwy, J. (2013). *The predictive mind*. Oxford: Oxford University Press.
- Kolmogorov, A. N. (1942). Determination of the center of scattering and the measure of accuracy by a limited number of observations. *Izvestiia Akademii Nauk SSSR. Series Mathematics*, 6, 3–32.
- Kolmogorov, A. N. & Hewitt, E. (1948). *Collection of articles on the theory of firing*. Santa Monica, CA: RAND Corporation.
- Kripke, S. (1980). *Naming and necessity*. Cambridge, MA: Harvard University Press.
- Lakoff, G. & Johnson, M. (1980). *Metaphors we live by*. Chicago: University of Chicago Press.
- (1999). *Philosophy in the flesh: The embodied mind and its challenge to Western thought*. New York: Basic Books.
- Loar, B. (1988). A new kind of content. In R. H. Grim & D. D. Merrill (Eds.) *Contents of thought: Proceedings of the Oberlin Colloquium in Philosophy*. (pp. 117–139). Tucson, Arizona, University of Arizona Press.
- Mackie, J. (1965). Causes and conditions. *American Philosophical Quarterly*, 2 (4), 245–264.
- O'Donovan-Anderson, M. (1997). *Content and comportment: On embodiment and the epistemic availability of the world*. Lanham, MD: Rowman & Littlefield.
- O'Regan, J. K. & Degenaar, J. (2014). Predictive processing, perceptual presence, and sensorimotor theory. *Cognitive Neuroscience*, 5 (2), 130–131. <https://dx.doi.org/10.1080/17588928.2014.907256>.
- Orlandi, N. (2014). *The innocent eye: Why vision is not a cognitive process*. New York: Oxford University Press.
- Oyama, S. (2000/1985). *The ontogeny of information: Developmental systems and evolution*. Durham, NC: Duke University Press.
- Segal, G. (2000). *A slim book about narrow content*. Cambridge, MA: MIT Press.
- Seth, A. K. (2014). A predictive processing theory of sensorimotor contingencies: Explaining the puzzle of perceptual presence and its absence in synesthesia. *Cognitive Neuroscience*, 5 (2), 97–118. <https://dx.doi.org/10.1080/17588928.2013.877880>.
- Titchener, E. B. (1929). *Systematic psychology: Prolegomena*. New York: MacMillan.
- Warren, W. (2005). Direct perception: The view from here. *Philosophical Topics*, 33 (1), 335–361.
- Wiese, W. (2016). What are the contents of representations in predictive processing? *Phenomenology and the Cognitive Sciences*, 1–22. <https://dx.doi.org/10.1007/s11097-016-9472-0>.
- Wiese, W. & Metzinger, T. (2017). Vanilla PP for philosophers: A primer on predictive processing. In T. Metzinger & W. Wiese (Eds.) *Philosophy and predictive processing*. Frankfurt am Main: MIND Group.