

---

# The Evidence of the Senses

## A Predictive Processing-Based Take on the Sellarsian Dilemma

Paweł Gładziejewski

---

Traditional foundationalist empiricist projects in epistemology postulated that sensory states of the subject are epistemically basic, in that they are capable of conferring justification on mental representations of the world without themselves needing to be (inferentially) justified by any antecedent representational states. This sort of view faces a seemingly hopeless dilemma, whose recognition is usually attributed to Wilfrid Sellars. If we treat sensory states as brute stimulations devoid of intentional content, then it is hard to see how the senses could provide subjects with anything that could possibly feature in justification-conferring relations with representational states. If we treat them as contentful, then in order to justify contentful states, sensory states themselves would presumably need to be justified by other representational states; but if this is so, they are not able to play a properly foundational epistemic role.

In the article, I use the Predictive Processing (PP) view of perception in order to sketch a possible resolution of the Sellarsian dilemma. I draw on PP in order to show how sensory states could actually serve a normative role that is recognizably similar to the one envisioned by traditional empiricists. To do this, I first distinguish representational from non-representational posits of PP and subsequently focus on the role that PP ascribes to sensory or “driving” signal. In particular, I argue that (1) the driving signal plays a role of a non-representational, contentless detector; at the same time, (2) it serves as an “impartial” or “theory-neutral” tribunal against which contentful internal models are actively tested and updated. Drawing on Anil Gupta’s work, I discuss the epistemic involvement of the sensory signal in perceptual inference and show how the signal provides conditional justification (i.e. justification that is conditional on the justification or rationality of prior knowledge) to perceptual hypotheses. Then I discuss the role the sensory signal plays in perceptual learning. I employ the notion of “epistemic convergence” to sketch out how the sensory signal could provide perceivers with unconditional justification (i.e. one that is not relativized to the justification of prior knowledge). If this approach is right, the Sellarsian dilemma seems to be averted. We can see how the senses can be at the same time silent (i.e. contentless) and capable of playing a sort of foundational epistemic role.

### 1 Introduction

One of the philosophically important aspects of the predictive processing framework (PP) is its postulate that our perceptual contact with the world is active and interpretative in nature. According to PP, perception relies on an internal generative model which estimates the states of the environment and on this basis predicts, in a top-down manner, the sensory states of the perceiver. The emphasis on the active role played in perception by the perceiver’s cognitive apparatus and prior beliefs is by no means historically new. But at the very least it constitutes a notable divergence from a view that has, in recent decades, been quite popular in naturalistically oriented philosophy of mind — a view on which perception (and perhaps representation more generally) is largely a matter of *detecting* states of affairs by being causally attuned to them (see Ramsey 2007, Chapter 4 for a discussion of the “receptor” notion of representation and its prevalence in philosophy and cognitive science).

### Keywords

Anil Gupta | Epistemic justification | Immanuel Kant | Kantian receptive sensibility | Perceptual inference | Perceptual justification | Predictive processing | Sellarsian dilemma

### Acknowledgments

Work on this paper was supported by the Polish National Science Centre FUGA 3 grant (UMO-2014/12/S/HS1/00343). I thank Michael Anderson, Jakob Hohwy and other participants of the Philosophy of Predictive Processing workshop at the Frankfurt Institute for Advanced Studies for their comments and questions. I am especially indebted to Jona Vance, Krys Dołęga, Thomas Metzinger and Wanja Wiese for their critical assessments of the earlier versions of the paper. When working on the paper, I also benefited from conversations with Marcin Miłkowski, Paweł Grabarczyk and Przemysław Nowakowski.

It is not surprising, then, that the view of perception as active in this sense caught the attention of philosophers (Clark 2013a; Clark 2013b; Clark 2016; Hohwy 2013). But perception is not *just* active according to PP. In addition to postulating the internal probabilistic inference that generates top-down predictions, the theory also attributes a crucial role to the sensory signal which results from the world's interaction with the perceiver's sensorium. It is through confronting the predicted with *actual* sensory states that the prediction error signal is generated and propagated bottom-up, which in turn enables the model to adjust its estimates. So despite having a clearly constructive flavor, PP rather views perception as resulting from the interplay between — to use Immanuel Kant's terminology — mind's spontaneity (i.e. the endogenous, prior-belief-based activity of the generative model) and receptive sensibility (i.e. sensory signal caused by the environment). In the present paper, I want to investigate more closely this latter, "receptive" aspect of PP, with emphasis on epistemological considerations.

Notably, traditional empiricist projects in philosophy used to attribute to the sensory states a foundational role. The sensory states were to act as epistemological "unmoved movers", capable of conferring justification or warrant on the subject's representational states, without needing a separate source of justification or warrant for themselves (BonJour 1985; Chisholm 1977). This empiricist-foundationalist project received deep and influential criticisms in the latter half of 20<sup>th</sup> century, to the point where it became often regarded as conclusively refuted (Davidson 1973; Quine 1969; Sellars 1956). One way to see what is wrong with this project is to consider a seemingly hopeless dilemma whose recognition is sometimes attributed to Wilfrid Sellars (Sellars 1956). If we treat sensory states as devoid of intentional content, then it is hard to see how they could possibly feature in inferential, justification-conferring relations with representational states. But if we treat sensory states as contentful, then in order to play a justificational role, they would presumably require to be justified by other, antecedent representational states. However, if this is so, then sensory states are not able to play a properly *foundational* epistemic role.

The aim of this paper is to revisit the Sellarsian dilemma with the PP account of perception in view. I will argue that although the sensory signal that PP appeals to is best construed as nonrepresentational, it nonetheless plays a crucial epistemic role. In particular, it serves as an "impartial" tribunal against which contentful internal models are actively tested and updated. This testing-and-updating is a probabilistically rational process, usefully interpreted as a form of abductive inference. If this is right, then we can see how the senses can be at the same time "silent" (i.e. contentless) and capable of playing a crucial role in providing justification<sup>1</sup> for representational states. The resulting view is distinct from classic empiricist foundationalism in many important respects, but it preserves the basic empiricist notion of "the evidence of the senses".

The discussion is structured as follows. In section 2, I discuss the Sellarsian dilemma in more detail. Next, in section 3, I provide a philosophical interpretation of PP which distinguishes representational from nonrepresentational posits of the theory. There, I also argue that the sensory signal belongs to the latter category. In section 4, I explain how, despite being nonrepresentational, the sensory signal still plays a major epistemic role according to PP, a role which is not completely unlike the role played by sensory "given" in traditional empiricism. In particular, I examine the epistemic involvement of the sensory signal in perceptual inference and perceptual learning.

1 Throughout this paper, I use the notions of "justification" and "warrant" to signify positive epistemic status very generally understood. For a representational state to be epistemically justified in this sense, it need not be produced by a subject that has a conscious or reflective grasp of the justification in question, or by way of conscious, intentionally controlled inference (for a similar approach, see e.g. the notion of "epistemic entitlement" as introduced in Burge 2003). On this encompassing approach, states that are formed through unconscious inferential processes — such as those postulated in PP — can count as epistemically evaluable.

## 2 Senses, their (Purported) Normative Role, and the Sellarsian Dilemma

It is easy to see how the senses could causally shape our representations of the world. On the one hand, the senses establish a world-mind causal link, as they are reliably affected by external stimuli. In fact, Kant is thought to have generalized this point into a claim that the very nature of sensibility lies in its constituting a passive capacity to be affected by things (Langton 1998; I will turn to this notion in section 4). On the other hand, the senses also establish a mind-mind causal link in that the (externally generated) sensory states are causally involved in producing subject's intentional states, such as perceptual beliefs.

It is philosophically crucial to distinguish this causal role of sensory states from the *normative* role that they play, at least according to classical, foundationalist versions of empiricism. This latter role is connected to a specific place that the senses supposedly occupy in the structure of justification or warrant of contentful states (BonJour 1985). Broadly speaking, the point is that the sensory contact with the world provides subjects with *reasons* for holding a belief or even a specific conception of the world (be it folk, scientific, or both). And these reasons are supposed to be epistemologically basic. Although many of our beliefs may acquire their justification by virtue of their inferential relations to other justified beliefs, this inferential chain ends with the deliverances of the senses. Through the sensory contact with the world, the subjects are supposed to be “immediately apprehending” states of affairs, or be “directly aware” of them, or be “acquainted” with them, or have those states of affairs “given”. Whatever the terminology, these sensory acquaintances can transmit justification or warrant to intentional states, but themselves do not require to be justified or warranted by being inferentially related to other contentful states.

Since serving as a reason for an intentional state is not the same as serving as a cause of that state, it seems impossible to account for the normative, justification-generating role of the senses by simply pointing to their causal role. In fact, as already mentioned, the very claim that we can even *make sense* of the normative role of sensory states faces the “Sellarsian” dilemma (Sellars 1956; see also BonJour 1985; Lyons 2008). Sensory states can be either taken as devoid of intentional content, or they can be treated as states that are contentful themselves. Both options seem to render them unfit to play their postulated normative role.

Accepting the first option amounts to accepting the claim that sensory states are contentless. They do not represent anything as being a certain way; rather, they constitute brute sensory stimulations. But brute sensory stimulations do not seem like things that could justify or warrant intentional states (Davidson 1986). In particular, they cannot enter inferential relations with intentional states. Because of this, sensory states, nonintentionally construed, are epistemically inert.

If we rather go with the second option and treat sensory states as contentful, then another problem arises. For any sensory state which asserts some content (i.e. represents the world as being in a particular way), we may legitimately ask about what justifies this state having (asserting) this particular content and not some other. After all, we presumably do not want this content to be arbitrarily asserted. Thus, if sensory states are contentful, then perhaps they can act as reasons, but they will also *require* reasons for themselves.<sup>2</sup> However, if this is the case, then the chain of inferential justification cannot end at the senses, hence, the latter cannot play a *foundational* epistemic role. Perhaps sensory representations do obtain their justification inferentially, from other, antecedently held representational states. But if so, then they are not “theory-neutral”. Their justification rests on conceptions and beliefs

<sup>2</sup> As one of the reviewers of this paper points out, an epistemologist might accept that sensory states are contentful (i.e. assessable for accuracy) and yet deny that they are rationally evaluable (i.e. assessable for their epistemic status). But note that although this is *some* way out of the Sellarsian dilemma, it does not seem to take us far. At the heart of the dilemma is the worry about how sensory states could provide reasons for beliefs. If a state is contentful but not rationally evaluable, then — as I understand it — we can establish what its conditions of accuracy are but not whether this state is justified or not (or whether the subject is justified or not in being in the state with that content). If so, it is not easy to see how the sensory state in question could generate a positive epistemic status for other representational states, i.e. play a justificational role. What we get is a contentful, but epistemically inert state.

which are merely presupposed (if we want to treat those background presuppositions as *empirically* justified, then the problem reappears, and so a regress looms). What the senses deliver is not bare, inferentially uncontaminated “Given”.

One possible reaction to the Sellarsian dilemma might be to claim that it rests on an overly restrictive, internalist and doxastic conception of epistemic justification, one captured in the dictum that “nothing can count as a reason for holding a belief except another belief” (Davidson 1986, p. 310). To pose the dilemma, one needs to assume that justification or warrant is a matter of inferential relations between intentional states of the subject. But perhaps the whole ordeal can be easily averted if we extend our notion of justification so that it counts, as possible justifiers, extra-mental and non-intentional factors, e.g. causal processes that reliably produce true or accurate representations (see Lyons 2008). With such a view in hand, it might be argued that the basic sensory representations do not require *inferential* justification; them having an appropriate, reliable *causal* history will suffice.

It is not my intention here to take stance on the internalism-externalism debate about epistemic justification. However, I do think that the Sellarsian dilemma *is* of relevance at least in the context of PP with its basic theoretical commitments. First, according to PP, perceivers estimate the states of the environment using only the resources that do not go beyond the statistical patterns that arise at their sensorium (Clark 2013b; Clark 2016; Hohwy 2013). In other words, perceivers update their model of the environment based solely on evidence that is internal to their mental lives. Second, PP construes this model updating in terms of unconscious probabilistic *inference* that employs prior (“presupposed”) representations of the causal structure of the environment. So the epistemology of perception inherent in PP seems distinctly internalist<sup>3</sup> and inferentialist. And these ideas are precisely the philosophical incubators of the Sellarsian dilemma. Indeed, if the present paper is on the right track, then it turns out that PP enables us to tackle the Sellarsian dilemma largely *on its own epistemological terms*, without relying on the externalist view of justification.

### 3 Representational and Nonrepresentational Posits of Predictive Processing

One of the philosophical discussions sparked by PP centers around the question of whether PP is best seen as representationalist, and if so, then what sort of internal representations it is committed to, exactly (Hohwy 2013; Hohwy forthcoming; Clark 2016; Gładziejewski 2016; Orlandi forthcoming). However, we need to be cautious to notice that it might be too simplistic to construe the debate as a competition between thoroughly representationalist and thoroughly antirepresentationalist readings of PP. To explain perception, PP postulates a complex processing scheme engaged in minimizing the prediction error. It is entirely possible that this scheme includes both representational and nonrepresentational aspects or parts. If so, then in order to get a complete philosophical understanding of PP, we need to carefully examine which of its theoretical posits belong to which category.

Following William Ramsey (2007), I take it that any genuinely representational theoretical posit of a cognitive theory owes its status to its meeting an appropriate “job description”, i.e. to the function that it plays in a cognitive or computational economy of a larger system. In other words, a theoretical posit counts as truly representational if it can be shown to *serve* as a representation in some nontrivial way or sense.

Now, we may take PP to be committed to at least following four posits: (1) the sensory signal which results from the world affecting the sensory apparatus of an organism, (2) the (hierarchical) generative model that sends top-down sensory predictions or “mock” sensory signals, (3) the prediction error signal which is propagated bottom-up and signifies the divergence between the predicted and actual sensory signal, (4) precision estimators which regulate the gain on prediction error signal. For each

<sup>3</sup> Because these epistemically relevant mental factors may be not consciously or reflexively accessible for the subject, the brand of epistemic internalism at play here is so-called “mentalism” rather than access internalism (see Conee and Feldman 2001).

of those posits, we may ask whether its functioning merits a representational reading. Here, I want to raise this question for (1) and (2) in particular, as they are directly relevant to the present discussion.

Let me start with posit (2), i.e. the generative model. There are strong reasons to regard it as playing a nontrivially representational role. First, it generates, in perceptual inference, estimates of the environment which guide the cognitive system's practical engagements with the environment (in active inference). Second, the model's ability to play this action-guiding function is dependent on how well its structure captures or resembles the causal structure — hierarchically nested at different timescales — of the environment. Third, the model does not simply passively register external states of affairs, but rather performs a sort of, largely endogenously-controlled, predictive simulation. It displays at least some degree of detachment or independence from current states of the world; in fact, a case could be even made that it can be used purely off-line, i.e. outside of any direct engagements with the environment. Fourth, insofar as the model undergoes correction in light of prediction error signals, it can be said to be capable of detecting cases when its estimates or hypotheses are inaccurate. The upshot, then, is that the generative model constitutes an action-guiding, detachable structural representation, capable of detecting representational error. The way this representation functions is not unlike the way that familiar and noncontroversial examples of external representational artifacts, like cartographic maps, function (for a much more in-depth exposition of this view, see [Gładziejewski 2016](#)).

If the generative model is in the business of representing the causal structure of the environment, then, according to PP, this representation is formed on the basis of the trace that the causal structure in question leaves on the outer boundary of the central nervous system. The job of the internal model is to recover the causal furnishing of the world using its sensory effects, i.e. the sensory signal. The question, now, is whether the sensory signal itself can be reasonably construed as representational.

Some might feel inclined to defend a representational reading of the sensory signal by pointing to the role of the senses as *detectors*. A detector is a structure that causally co-varies with certain states of affairs, and its function is to do so. In particular, detector's function is to reliably react to some conditions and in turn initiate certain reactions or downstream effects. Unsurprisingly, the sensory signal can be regarded as detecting the presence of environmental conditions. Sensory signals result from a reliable causal covariance between the sensory apparatus and the presence of proximal physical stimuli. And given the larger processing scheme that they participate in (according to PP), it is the function of the senses to be reliably affected by the stimuli. In particular, they give rise to prediction error signal, as unpredicted or “unexplained” aspects of the sensory signal get propagated bottom-up, eventually prompting the generative model to revise its estimates of the environment, if needs be.

However, it is far from clear how functioning as a detector could be by itself *sufficient* to regard something as representational ([Orlandi forthcoming](#); [Ramsey 2007](#)). This is because any causal mediator meets the conditions of serving as a detector. On such a view, a light switch, a firing pin in a pistol, or the reactions of an immune system to bodily injury would have to count as representational. Claiming that the senses “represent” anything in this extremely liberal sense would amount to making an uninteresting claim that they causally mediate between what happens in the world and what happens in the perceiver (i.e. that they are transducers). In other words, attributing the sensory signal with a role of representing something simply because it detects proximal conditions would mean subscribing to a hopelessly trivial notion of representation.

The case for a representational reading of the driving signal weakens even further when we notice that it lacks other functional properties that we often attribute to representations (see also [Orlandi forthcoming](#)). First, the signal is not directly involved in guiding the actions (active inferences) of the cognitive system. What guides active inference is not the sensory input itself, but rather the interpretation of the input provided by the generative model in the form of perceptual hypotheses. Second, signals generated in the sensorium necessarily track current proximal stimuli; they only react to what is actually present. They do not entertain any sort of independence from their physical causes. As such, the sensory signal is incapable of acting in a way that is even minimally detached or off-line.

Given how there is no representation without the possibility of representational error, it is also natural to consider the question of whether we could make sense of the claim that the incoming sensory signals are capable of *misrepresentation*. Admittedly, it makes perfect sense within the PP framework to say that the senses can sometimes mislead or misguide the perceiver, say when she tries to find her way in a deep fog. These are the cases where the sensory signal is noisy — in the sense of having low precision or inverse variance — to the point where it tends to lead perceptual inference astray. But it does not seem quite right to count such cases as examples of the *senses* somehow failing at performing their duties, let alone misrepresenting anything. The crucial point to note here is illuminatingly expressed by Anil Gupta:

If, during a walk in a forest, I bump my head on a low branch of a tree, it is better that I assume responsibility (and change my ways) than that I pin the blame on the tree. The tree is passive. It is bound to be the way it is, given the circumstances, and it is useless to blame it for my sore head. Similarly, if, having suffered an experience, I acquire a false perceptual belief, it is better that I assume responsibility (and change my manner of “reading” experience) than that I pin the blame on the experience. The experience is bound to be the way it is, given the circumstances, and it is useless to blame it for my false belief. (Gupta 2006, pp. 28–29)

This idea can be in the following way translated into more subpersonal notions that PP trades in. When we qualify noisy sensory signal as misleading, unreliable or ambiguous, it is not due to the fact that the sensory apparatus somehow fails to represent the world correctly. Photoreceptors in the retina do not fail to fulfill their function, representational or otherwise, when we are taking a stroll on a foggy day. They are “bound” to generate a noisy reaction, given the circumstances, and it is useless to “blame” them, i.e. treat as malfunctioning<sup>4</sup>. Rather, the qualification of the sensory signal as misleading is parasitic on the use that the inferential machinery, i.e. the generative model, makes of the signal.<sup>5</sup> It is the model’s job to provide hypotheses about determinate causal etiology of the incoming signal. The noisy sensory signal itself is misleading only in the sense that it can make the generative model more prone to come up with inaccurate hypotheses about the signal’s causal origins, and hence less effective at minimizing the prediction error. It is not the signal itself that misrepresent the world, but the model that “misreads” the signal and misrepresents its distal causal etiology.

These considerations present a case to think that the explanatory repository of PP involves both representational and nonrepresentational posits. In particular, while the generative model can be treated as playing a genuinely representational role, the same cannot be said about the sensory signal. What the *senses* deliver, then, is *not* representations. This echoes Immanuel Kant’s claim that “although it is correct to say that the senses do not err, this is not so because they always judge correctly but because they do not judge at all” (Kant 1781/1996, p. 128).

## 4 The Sellarsian Dilemma in Light of Predictive Processing

### 4.1 Epistemic Role of the Senses According to Predictive Processing: an Outline

The foregoing discussion suggests that if we are to tackle the Sellarsian dilemma using PP’s conceptual resources, we will have to proceed by accepting the claim that the sensory signal can somehow play a sort of foundational epistemic role despite the fact that it is not contentful. However, if we assume,

- 4 This is *not* to say that the sensory apparatus cannot malfunction in cases where the physiology of the sensory organ itself is damaged or otherwise changed. However, this sort of malfunctioning is *not* an instance of *misrepresentation* (assuming that after the damage or modification, the sensory organ still functions as a causal mediator, and so does not have representational content).
- 5 A remark is in order here. When an inaccurate representation is formed based on a noisy signal, there is usually not one, but two culprits: the model which misestimates the state of the environment and the precision estimators which fail to lower (to a sufficient degree) the gain on prediction error signals.

rather uncontroversially, that inferences are transitions between *representational* states which accord to some normative epistemic rule, then contentless sensory signals turn out incapable of acting as contentful *premises* from which the latter could be inferentially derived. The question, then, is how to characterize the epistemic or rational, and recognizably “foundational” bearing that the sensory signal has — assuming that it has — on perceiver’s internal representations?

To avoid confusion, we need to respect the distinction between sensation and perception. In PP, sensation corresponds to the contentless sensory signal.<sup>6</sup> Perception, construed in terms of perceptual inference, corresponds to the forming of a representation of the sensory signal’s distal causal origins. One might wonder whether we should rather go with the second option of the dilemma and treat the perceptual representations as epistemically basic. But it seems that perceptual representations or hypotheses, as PP construes them, do not even purport to be epistemically basic. According to PP, cognitive systems form perceptual hypotheses by engaging in approximate Bayesian inference which relies on preexisting probabilistic representations of (or prior beliefs about) the causal structure of the environment (see section 4.2). Thus, perception rests on presuppositions and is inherently theory-laden in this sense. What cognitive systems perceive is inferentially shaped by what they already represent the world to be like.<sup>7</sup>

But this is not the whole story about how perceptual representations are formed in PP. Insofar as perception is prediction error minimization, updating perceptual representations relies not *only* on preexisting beliefs but also depends on the incoming sensory signal. After all, the prediction error is estimated and minimized relative to actual sensory input. Of course, we need to be careful to distinguish between sensory signal’s causal involvement in forming representations and its epistemically normative involvement. The present point is that the sensory signal can be attributed with not only a causal but also a normative, epistemic role.

Remember the Kantian lesson (which also resonates in the passage from Gupta cited in the previous section) that our sensory contact with the world is *receptive* in the sense that it constitutes a passive capacity to be affected by things; the raw sensory manifold is simply received rather than endogenously constructed through the spontaneous activity of the perceiver (Langton 1998). The sensory input, as construed by PP, *is* receptive in this sense. The states of the sensory apparatus are treated as a function of their hidden worldly causes (Friston 2010; Friston and Kiebel 2009). What happens at the perceiver’s sensory boundary is outside her jurisdiction, in the sense of being contingent solely on the causal structure of the world (which includes the perceiver’s own body), and independent from the perceiver’s prior beliefs or inferential activity. Metaphorically, the senses constitute a point at which inference and top-down prediction bottoms out, and the system becomes purely responsive to external factors.

6 There is a way in which the present PP-based construal of sensory states diverges from epistemological tradition. Traditionally, the purportedly foundational sensory states were assumed to have a qualitative or phenomenal character. They were often seen as raw conscious “feels”. Here, they are treated as a bottom-up signal that the results from the world causally affecting the sensory apparatus of the perceiver. I do not assume that sensory states thus understood have to determine or be manifest in the phenomenal perceptual experience of the subject, even if they shape it in major ways. In the context of PP, it seems more reasonable to claim that contents of consciousness correspond to generative-model-derived perceptual hypotheses which already populate the world with familiar objects, properties and relations (see the discussion of sensation/perception distinction in main text). The analogy between “sensation” in the present sense and the way this notion was used in more traditional approaches *only* pertains to the epistemic role (see the following discussion in main text).

7 As has been already noted in the literature (Hohwy 2013; Lupyán 2015), the issue of cognitive penetrability of perception naturally crops up in the context of PP. Cognitive penetrability is sometimes construed in terms of, roughly, high-level cognitive states (which may correspond to personal-level beliefs or desires) directly affecting the content of perceptual representations. Cognitive penetrability thus understood is a local phenomenon, as many percepts seem *not* to be penetrable in this sense. For example, in the hollow mask illusion, the mask persists to be perceptually represented as convex even in light of knowledge that it is actually concave. But notice that PP commits us to the view that perception is inferential all the way down. Perceiving may be sometimes independent of representations stored at *higher* levels of the generative model, but it still rests on *some* assumptions, presumably stored at lower levels of the processing hierarchy. For example, on PP view of things, the supposedly convex shape of the hollow mask is inferred from preexisting premises that are low-level and beyond the conscious control of the subject. When I talk about perceptual representations being theory-laden in PP, I mean that they are strongly or globally cognitively penetrable in this sense. Notice that this strong penetrability of perception is presumably what we should care about in the context of foundationalist empiricism. When epistemologists seek properly epistemically basic representations, they do not mean representations whose justification does not depend on a particular subclass of intentional states (say, personal-level beliefs about the truth of particular scientific theories; see Fodor 1984), but ones whose justification does not depend on (inferential relations to) any antecedent intentional states.

By being passive in this sense, the sensory signal is anything but theory-laden. Rather, it is “pure” or “impartial”, as it only depends on what the worldly causes are, not on what they are represented to be.<sup>8</sup>

The claim about the receptivity of the sensory signal requires qualifications. There is a sense in which the sensory signal *does* depend on the activity of the perceiver and her representations of the environment. For one thing, perceptual inference is intertwined with active inference, whereby the perceiver actively samples the environment in order to *make* the sensory input fit the perceptual hypotheses. Perceivers intervene in the world with their bodies to cause it to conform with their sensory predictions. Also, through estimating precision of the sensory input, the perceiver has the ability to regulate the weight of the prediction error signal, and hence the degree to which it affects perceptual inference. For example, in a deep fog, the perceiver may “decide” to largely ignore the noisy visual input. However, on closer inspection, neither of those points is inconsistent with the claim about the essentially receptive nature of the sensory signal. *Once* the perceiver decided to actively sample the environment in a particular way (say, by performing a particular series of saccades) and *once* the precision of the input is estimated, it is no longer up to the perceiver what will actually happen in her sensorium; it depends on the external causes. Whether active inference actually succeeds at making incoming signal conform to the signal predicted requires “cooperation” on the part of the external environment. Similarly, although the perceiver may lower the degree to which the sensory input affects hypothesis revision, it does not thereby change the sensory signal itself; again, it depends on the world.

The claim about the receptivity of the senses relates not only to how they are causally situated, but it has a significant epistemological import as well. To get a general gist of how the epistemic role of the sensory signal is to be understood, imagine a cognitive system which lacks anything that fulfills this role. That is, imagine a system which gets sensorily disconnected from the world — say, we disable its sensory system — but still uses an internal generative model to get around. The system samples from its model (construed here as a Bayesian network) and thereby performs an internal predictive simulation of the causal processes in the world, on the basis of which it decides how to act. The simulation in question relies *only* on prior and likelihood probability distributions already encoded in the model. No sensory input is received, so no prediction error is computed. Our system freely “dreams” reality (see [Hobson and Friston 2012](#); [Bucci and Grasso 2017](#)),<sup>9</sup> rather than properly perceives it.

Even if we assume that the hypothesis (or a set of hypotheses) which initiates internal sampling from the model actually corresponds to some initial state of the environment, given the complexity and unpredictability of the latter, it is natural to expect that the internal simulation will at some point diverge from actual states of the world, eventually leading to potentially catastrophic results. Without the sensory input with which the internal predictions are confronted, the process of revising representations lacks some sort of *external constraint* on the space of currently relevant hypotheses, a constraint which makes the process in question responsive to what is actually going on in the environment. When the sensory constraint or guide of this sort is not present, the internal simulation is like, to use John McDowell’s (1994, p. 11) poetic phrase, “frictionless spinning in the void”. And this shows what the epistemic role of the senses consists in. As active and constructive as perception is according to PP, if it is to put us in touch with the world, there needs to be a point where it meets resistance or friction and becomes answerable to some external authority. This is precisely the job of the sensory signal with its passive or receptive nature. Perceivers minimize the prediction error which reliably depends on both internally-generated predictions and actual sensory input; but the latter ultimately

<sup>8</sup> This claim about “purity” of the senses needs to be qualified. The physical/physiological makeup of sensory transducers themselves is, of course, a product of evolutionary forces. The senses are selective with respect to the physical energies they react to, in a way that expresses general “assumptions” what the organism’s Umwelt is (roughly, about what in the environment is biologically useful or salient). It is not my intention here to treat the senses as providing some kind of “view from nowhere”. The senses count as passive receivers, but what they receive (and how they receive it) is determined by the evolutionary history. I am indebted to Thomas Metzinger for pushing me on this.

<sup>9</sup> As one reviewer points out, this analogy with dreaming is limited in two important ways. First, certain kinds of sensory signals (e.g. proprioceptive and interoceptive) are actually present during sleep. Second, the content of perceptual-like states in dreaming is less detailed than the content of actual perceptual states.

depends on worldly causes. This way, through in the involvement of the sensory input, perception becomes supervised by the world itself (Hohwy 2013).<sup>10</sup>

Crucially, for the senses to play this role, they do not have to supply perceivers with epistemically basic *representations*. The sensory apparatus does not provide prediction-error-minimizing systems with a set of ultimate, non-inferentially justified premises, but rather with raw, uninterpreted data against which sensory predictions are tested. The system can then use the result of this testing — the prediction error signal — to correct its internal representation. The causal structure that generates the sensory input is *indirectly* recovered by adjusting the model to optimize prediction error minimization. It is not simply inferred from a set of ready-made, pristine premises, “direct acquaintances”, or anything of that sort.

This provides us with a general understating of what the “evidence of the senses” amounts to, according to PP. Now I want to fill this general story with some details, by discussing the epistemic role that the sensory signal plays in perceptual inference (section 4.2) and perceptual learning (section 4.3).

## 4.2 Sensory Signal and the Epistemology of Perceptual Inference

In a way, epistemological considerations lie at the very heart of PP. Despite being naturalistic through and through, the theory views perception as a process that conforms to *normative* principles. This means that, on PP view of things, organisms in fact (tend to) form perceptual representations that they *ought* to form *were* they to follow a rational norm (Hohwy 2013; Hohwy et al. 2008; see also a related discussion in Rescorla 2016). One way to understand the norm in question is in terms of causal abduction (inference to the best explanation), whereby percepts are hypotheses that “best explain” the sensory input by citing its worldly cause(s) as an explanans (Hohwy 2014; Seth 2015). As already mentioned, in PP, this general idea is cashed out in Bayesian terms. The assumption is that perceivers update (infer) perceptual hypotheses in a way that aims at maximizing posterior probability. This process is only approximately Bayesian, as PP is not committed to the claim that brain literally works by implementing Bayes’ rule (Hohwy 2013). Rather, the point is that the process of prediction error minimization with the use of generative models approximates Bayesian inference. A system updating its generative model to minimize prediction error is a system that updates its internal estimates of the environment in a way that conforms with Bayes’ rule. The upshot, then, is that PP construes perception as a form of probabilistically rational abduction.<sup>11</sup>

Given that perceptual inference is rational in this sense, how are we to understand the epistemic involvement of sensory signal in this process? The preceding section already characterized this role in broad strokes. To a first approximation, sensory input constrains “from the outside” the updating of hypotheses in a way that enables them to actually track environmental states of affairs. Now I want to

<sup>10</sup> The claim about the world acting as a supervisor does not imply that the notion of justification at use here is externalist after all. The process of internal model formation (and perceptual hypothesis formation) is truth-conducive in the sense that if the sensory signal one receives is actually causally generated by the familiar physical world, then one will likely form — on the basis of the input — a more or less accurate representation of the world. This does not have to imply that an actual reliable causal connection to the world acts as a justifier in this story. To understand this point, consider a PP-inspired variation on the so-called “new evil demon” scenario (Cohen 1984). Imagine your epistemic copy that, through life, receives a series of sensory signals which are exactly like the signals that you receive. Imagine that this copy performs a series of (subpersonal) probabilistic inferences to construct a model of the worldly causes of its sensory states — again, the copy proceeds exactly like you do (or your brain does). The copy ends up with a model of the external world which is exactly like yours. Given the same sensory signals, it will form the same perceptual hypotheses. Now imagine that this copy is actually fed its sensory signals by an evil, misleading demon, such that the copy has no reliable causal connection to a real world. The model of the world harbored by your copy is systematically false. But it is constructed by way of the same Bayesian inferences operating on the same data as is the case with you. So it seems that your copy is, given all available sensory evidence, as rationally entitled to/justified in forming the model of the environment as you are.

<sup>11</sup> I take it that this inferential aspect is inherent in Bayesian accounts of perception, including PP. It is hard to see how perception can be truly Bayesian without constituting Bayesian inference (see Hohwy forthcoming; Kiefer 2017; Rescorla 2015). But it needs to be acknowledged that some authors make attempts to reconstrue PP in a way that aims to avoid this inferential commitment (see Orlandi forthcoming). It is beyond the scope of the present paper to discuss how this anti-inferential reconstrual could affect the issues about the epistemology of perception in the context of PP.

propose how we can enrich this story by linking it to some ideas found in Anil Gupta's "Empiricism and Experience" (Gupta 2006).

Gupta's aim is to understand the "rational contribution" that conscious experience of the subject makes to knowledge. Right now I undergo a certain visual experience which, intuitively speaking, makes *reasonable* my perceptual belief that there is a computer screen in front of me. Gupta coins this rational contribution of my experience to my belief as "the given". On his account, the given in any experience is not, so to speak, epistemically autonomous. In his words, every experience is "multiply factorizable" in the sense that "no experience carries with it its own genealogy" (Gupta 2006, p. 7). For example, an experience that one has when looking at a bright green wall so large that it occupies one's whole visual field could have been obtained by looking at a white wall through bright green glasses, or at a blue wall through yellow glasses. Experience *alone* cannot decide between these options. This gives rise to the claim that the given in experience is "hypothetical" in nature. That which judgment is rational to hold on the basis of a particular experience is conditional on the background "view" that one brings to bear on that experience. The view consists of "concepts, conceptions and beliefs" that perceiver is already in possession of. For example, if I antecedently believe myself to be placed in a large building, and I am in the possession of the concept wall, and my view does not include the belief that I am wearing yellow glasses, etc., then, given a particular visual experience, I can reasonably form the perceptual belief that there is a bright green wall in front of me. A person with a different view could well be entitled to a different judgment, given the same experience. This leads Gupta to conclude that the "logical category" of the given in a particular experience  $e$  is that of a *function*  $\Gamma e$ , which takes the view  $v$  as the input and yields (classes of) perceptual judgments  $\Gamma e(v)$  as outputs. The resulting position can be summarized with the following schema:

View  $v \Rightarrow$  (Experience  $e \Rightarrow$  Perceptual judgments  $\Gamma e(v)$ )

On this view, then, experience provides the perceiver not with an absolute, but only with conditional entitlement to her perceptual judgment(s); one could say that perceptual judgments are "rationally underdetermined" by experience. The rationality (justification or warrant) of the perceptual judgment is conditional on the rationality (justification or warrant) of the view that one brings with herself. Importantly, the *whole* positive nature of the given in experience consists in how it provides rational guidance in forming perceptual judgments. In particular, although the given in experience outputs judgments, it is not itself contentful or "propositional".<sup>12</sup> Taken apart from the view, the experience as such is "silent", rather than composed of basic propositions regarding sense-data, subjective mental states or medium-sized physical objects.

Gupta construes his account as concerned with strictly logical or normative matters, and distances it from naturalism, claiming that "the project of constructing a naturalistic account of rationality and of perceptual judgment is, at the present stage of inquiry, nothing but quixotic" (Gupta 2006, p. 54). Nonetheless, against the author's intentions, I want to suggest that his treatment of "the given" proves strikingly useful as a framework for understanding the epistemic role that the sensory signal plays in perceptual inference according to the naturalistic theory that is PP.

For starters, we may observe that basic categories present in Gupta's account have analogues in PP's story. Namely, we may suppose that (1) the raw sensory signal corresponds to Gupta's "experience"<sup>13</sup>; (2) the generative model that encodes the causal structure of the environment corresponds to "the

<sup>12</sup> When Gupta 2006 characterizes the given as nonpropositional, the claim is not (only) that it lacks "conceptual" or "propositional" type of content, but that it lacks content *altogether*.

<sup>13</sup> Importantly, the analogy here is supposed to pertain only to the similarity at the level of the type of normative role that the analogues play. In Gupta's view, the given is manifest in conscious experience (Gupta 2006, p. 30). But if my proposal is right, this particular point of Gupta's view is not preserved in PP. To repeat, it is not my suggestion that conscious perceptual *experience* is determined by or somehow corresponds to what happens at the level the sensory input.

view”; (3) perceptual hypotheses (estimates) generated by the model to minimize the prediction error relative to the sensory input correspond to perceptual judgments entailed or yielded by the experience when combined with the view. Abstractly, we might treat the perceptual representation or hypothesis as an output of a function  $\Gamma_s$  (where  $s$  stands for sensory signal) that takes the generative model (i.e. a particular prior and likelihood distribution)  $m$  as input. We may then reformulate Gupta’s schema as follows:

Generative model  $m \Rightarrow$  (Sensory signal  $s$  Perceptual hypothesis(es)  $\Gamma_s(m)$ )

Paralleling Gupta’s account, the transition from the combination of the incoming signal and the preexisting model of the environment to the perceptual hypothesis is rational. Namely, it is rational, as it (approximately) conforms to Bayes’ rule; and Bayesian inference is by its nature truth-conductive: if you follow the rule, you produce hypotheses with the largest posterior probability. Also, on this picture, the normative contribution of the sensory signal consists entirely in the guidance that it provides for the generative model. Apart from the generative model that interprets it, the sensory input is silent (contentless). This is PP’s incarnation of Gupta’s claim that the given in experience is “nonpropositional”. Furthermore, the guidance that the senses provide for the model is essentially corrective in nature. Its epistemic value lies in the fact that it enables the cognitive system to be responsive to cases in which what actually happens in the environment diverges from what is already “contained” (encoded) in, and hence expected on the basis of, the model alone. This, after all, is the lesson to be drawn from the dreamer-perceiver example given earlier.

Notice that on such a view the epistemic role the sensory signal does *not* consist in the signal *transferring* its epistemic status to the perceptual hypothesis. The signal as such is not rationally evaluable. Its rational involvement consists solely in how it is essential (in combination with the internal model) in enabling perceptual inference to be truth-conductive in virtue of approximately embodying Bayesian rationality. Without the signal, the rational link between priors and the perceptual hypothesis is broken. That is, were perceptual inference not guided by the senses, it would fail to be Bayes-rational.<sup>14</sup>

Crucially, this picture parallels Gupta’s proposal in that, on PP view of things, perceivers obtain only conditional epistemic entitlement (justification or warrant) to their perceptual hypotheses. Imagine two perceivers facing the same ambiguous stimulus, like the duck-rabbit. If one of them ascribes a significantly larger prior probability to ducks than to rabbits, and the other vice versa, they will (all else things being equal) end up forming different perceptual representations. But both representations can be considered equally probabilistically rational in a conditional sense. In each case, the hypothesis is formed which has the largest posterior probability in light of the probabilistic generative model used to interpret the sensory input. Both perceivers can be said to have conditional entitlement to their perceptual hypotheses — they are entitled to their perceptual representations *if* they are right about their respective priors.

### 4.3 Sensory Signal and the Epistemology of Perceptual Learning

Can perception, as PP construes it, give rise to justification that is unconditional or not relativized to a particular set of background “beliefs” encoded in the generative model? This is a subtle issue, and it

<sup>14</sup> Here, it is important not to overplay the causality/normativity distinction. Although the distinction needs to be drawn to avoid basic philosophical confusions, the point is *not* that matters of causality are completely disconnected from normative matters. The transition from the signal-model combo to a perceptual hypothesis is rational in virtue of its conforming to a rational, truth-conductive principle. But in PP, the move from the signal-model combo to the hypothesis is, of course, *causally realized*. Furthermore, one could argue that the rationality of token perceptual inferences lies precisely in the fact that their causal profiles track or accord to a rule of inference (here, Bayes’ rule; see also Kiefer 2017). The sensory signal’s normative involvement is constituted by how it shapes the *causal* transitions between contentful states in such a way that they accord to a *normative* rule. Thanks to Thomas Metzinger and Jona Vance for pointing this issue out to me.

is not my aim here to settle it definitely. Nonetheless, let me tentatively sketch out how this problem could be handled in a way that suggests a positive answer to the original question.

A crucial observation to make is that although perceptual inference makes use of a pre-existing model of the environment, the model as such is subject to change through perceptual learning (Clark 2013a; Clark 2013b; Clark 2016). If we treat the model as a hypothesis space, then this space is anything but fixed once and for all. Even if perceptual inference is only as good as the range of available hypotheses, the latter can change and improve over time through learning. One of the factors that make learning possible is, again, the receptivity of the senses. Because the world is basically independent of how it is represented in one's generative model, the patterns that arise in the sensorium can, and do leave average residual prediction error. This way, the sensory input enables the system to “recognize” that the world escapes even the best predictions that its current model can afford. In learning, this fact is exploited to optimize the effectiveness of the model at minimizing the prediction error over longer periods of time. At heart, learning consists in modifying the model parameters, presumably with the use of some form of gradient descent algorithm (Clark 2013b; Friston 2003; Friston and Kiebel 2009). Starting from an initial “guess”, the hierarchical structure of the model is iteratively adjusted so as to improve its overall ability to minimize prediction error. In other words, not only perceptual inference (i.e. forming hypotheses to deal with current sensory input) but also perceptual learning (i.e. changing the structure of the model itself) is basically prediction error minimization, performed at different time scales (Clark 2013b).

Importantly, in learning thus construed, the causal structure of the environment is gradually recovered from the sensory input itself. This process does not require supervision by an external observer providing the cognitive system with data that are already pre-classified (Hohwy 2013; see also Eliasmith 2005). Because what is available at the outset is basically raw data (perhaps accompanied with relatively minimal learning biases), learning on this view can be seen as a form of bootstrapping (Clark 2013a; Clark 2013b; Clark 2016). Furthermore, a body of computational work suggests that systems which learn using multi-level Bayesian strategies are able to recover the values of hyperpriors or “overhypotheses” (see the work on learning with Hierarchical Bayesian Models: Kemp et al. 2007; Kemp and Tenenbaum 2008; Tenenbaum et al. 2011; a summary can be found in Clark 2016). These are high-level priors that structure the range of hypotheses available at lower levels of the processing hierarchy. This way the learner can recover “deep” or abstract overarching principles of how the data are to be categorized, like the shape bias in categorizing physical objects (see Clark 2016, pp. 171–175). Hyperpriors are themselves selected by maximizing posterior probability. This means that a hypothesis space available at a lower level is selected (out of a range of possible such spaces) based on whether it is most likely to capture the causal structure that generated the data (Kemp et al. 2007).

But what consequences do these considerations have for the question of unconditional justification or warrant, assuming that they do? Here is a way of thinking about this, albeit admittedly sketchy and speculative. Consider an approach that, in a way, puts classical empiricist foundationalism on its head. On this view, representations that are unconditionally justified are not the starting points of the inquiry but, so to speak, its *endpoints*. The crux of this proposal has been formulated by Charles Sanders Peirce:

Different minds may set out with the most antagonistic views, but the progress of investigation carries them by a force outside of themselves to one and the same conclusion. This activity of thought by which we are carried, not where we wish, but to a fore-ordained goal, is like the operation of destiny. No modification of the point of view taken, no selection of other facts for study, no natural bent of mind even, can enable a man to escape the predestinate opinion. This great law is embodied in the conception of truth and reality. (Peirce 1878/2011, p. 63)

On this general approach, the notion of unconditional justification (or rational entitlement) is cashed out in terms of what we might call “epistemic convergence”. Roughly, the idea is that there are beliefs that an agent will eventually form, regardless of what she accepts at the starting point, if she proceeds by rationally updating her beliefs in light of evidence. These are the beliefs on which rational inquiry converges, and they are justified in an unconditional sense. Gupta, who himself adopts this approach (Gupta 2006), expresses this in terms of a convergence that results from a revision sequence, in which successive experiences  $\langle e_0, e_1, e_2, \dots, e_n, \dots \rangle$  result in a sequence of views  $\langle v_1, v_2, v_3, \dots, v_n, \dots \rangle$ , where each previous experience revises subsequent view. Formal and philosophical technicalities aside, the idea is that revision sequences may converge on particular views or propositions, regardless of their starting points. For example, a proposition  $p$  constitutes a point of convergence at stage  $n$  of the revision sequence generated by a particular series of experiences  $E$  if, and only if, regardless of the starting point  $v_0$ , all views that survive revision in light of  $E$  include  $p$  at  $m \geq n$ . A proposition on which the revision process eventually converges regardless of where it starts is a proposition which the subject is rationally entitled to accept in an unconditional sense.

It seems that the general notion of epistemic convergence can be accommodated by PP. The hypothesis is that systems which learn by iteratively revising their generative models to optimize long-term prediction error minimization can, and perhaps do reach points of convergence. And we may treat the representations on which such systems converge, regardless particular starting points and learning trajectories, as ones which are justified in an unconditional sense. For example, we may speculate that cognitively healthy people (as well as, perhaps, members of many other species) inhabit a perceptual world that is roughly the same in terms of its basic ontological furnishing. It is a world filled with medium-sized physical objects, which persist through time, occupy spatial locations, can retain their identities despite changing their properties, and which undergo causal interactions with each other. From the perspective of PP, this unity of how the world is perceptually represented may stem precisely from a sort of epistemic convergence. Prediction error minimizing systems (with relevantly similar bodies and sensory apparati) will converge to acquire (learn) generative models that parse the world into ordinary physical objects. This convergence point may be actually reached quite early in ontogeny. But it still presumably results from a bootstrapping process in which the values of hyperpriors are adjusted to constrain the range of hypotheses available at lower levels. The common-sense perceptual ontology could then be treated as unconditionally justified in virtue of it constituting such a convergence point of iterative model revision. Crucially for the present discussion, the “driving” influence of the sensory signal has major significance in enabling such convergence to occur. The convergence is possible because the sensory signals that organisms receive embody similar statistical regularities (which is in turn due to the signals actually being generated by the same physical world), so that the learning processes in different agents tend to proceed, at least at some level, in uniform direction. In other words, it is the sensory signal that constitutes the “outside force” which inevitably carries the learning process to “one and the same conclusion”.

## 5 Conclusion

The basic tenet of empiricism is that we owe (a large chunk of) our knowledge to the epistemic authority of “the evidence of the senses”. When we perceive, our thinking becomes answerable to the empirical world. Despite intuitive plausibility, this claim proves notoriously difficult to defend or even formulate in a way that is not philosophically problematic. One problem that empiricism needs to face is the Sellarsian dilemma, which purports to show that sensory states are (if we consider them to be nonintentional) incapable of conferring justification on intentional states, or (if we construe them as intentional) incapable of terminating the chain of inferential justifications.

In the present article, I attempted to reconsider the Sellarsian dilemma from the perspective of the Predictive Processing (PP) view of perception. According to PP, perception centrally involves an active

or constructive aspect, whereby an internal generative model attempts to predict the inflow of sensory signals. But it also involves an aspect that is “receptive” in roughly Kantian sense, as the model adjusts its predictions to match the signal that results from the world impinging on the perceiver’s sensory apparatus. On such a view, the epistemic authority of the senses lies in how the sensory signal constrains the activity of the model, making it “answerable” (through error-correction) to external states of affairs. The justification or rational entitlement that perceivers thus obtain to their perceptual hypotheses is conditional, as it depends on the epistemic standing of the generative model that interprets the incoming signal. However, because of how, in perceptual learning, the model itself is corrected using the guidance of sensory input, we can at least begin to make sense of the idea that our sensory contact with the world can provide us with unconditional justification or entitlement as well. The resulting view significantly diverges from classical empiricist foundationalisms which saw the evidence of the senses as consisting in basic, non-inferentially justified representational states. Contrary to this tradition, PP rather leads us to view the sensory signal as a purely contentless “foundation”, one that passively causally registers, rather than represents, what happens in the empirical world.

## References

- BonJour, L. (1985). *The structure of empirical knowledge*. Harvard: Harvard University Press.
- Bucci, A. & Grasso, M. (2017). Sleep and dreaming in the predictive processing framework. In T. Metzinger & W. Wiese (Eds.) *Philosophy and predictive processing*. Frankfurt am Main: MIND-Group.
- Burge, T. (2003). Perceptual entitlement. *Philosophy and Phenomenological Research*, LXVII, 503-548.
- Chisholm, R. (1977). *Theory of knowledge*. Englewood Cliffs, NJ: Prentice-Hall.
- Clark, A. (2013a). Expecting the world: Perception, prediction and the origins of human knowledge. *The Journal of Philosophy*, CX, 469-496.
- (2013b). Whatever next? Predictive brains, situated agents and the future of cognitive science. *Behavioral and Brain Sciences*, 36, 181-204.
- (2016). *Surfing uncertainty. Prediction, action, and the embodied mind*. Oxford: Oxford University Press.
- Cohen, S. (1984). Justification and truth. *Philosophical Studies*, 46, 279-295.
- Conee, E. & Feldman, R. (2001). Internalism defended. *American Philosophical Quarterly*, 38, 1-18.
- Davidson, D. (1973). On the very idea of a conceptual scheme. *Proceedings and Addresses of the American Philosophical Association*, 47, 5-20.
- (1986). A coherence theory of truth and knowledge. In E. Lepore (Ed.) *Truth and interpretation: Perspectives on the philosophy of Donald Davidson* (pp. 307-319). Oxford: Blackwell.
- Eliasmith, C. (2005). A new perspective on representational problems. *Journal of Cognitive Science*, 6, 97-123.
- Fodor, J. (1984). Observation reconsidered. *Philosophy of Science*, 51, 23-43.
- Friston, K. J. (2003). Learning and inference in the brain. *Neural Networks*, 16, 1325-1352.
- (2010). The free-energy principle: A unified brain theory? *Nature Neuroscience*, 11, 127-138.
- Friston, K. J. & Kiebel, S. (2009). Predictive coding under the free-energy principle. *Philosophical Transactions of Royal Society B*, 364, 1211-1221.
- Gupta, A. (2006). *Empiricism and experience*. Oxford: Oxford University Press.
- Gładziejewski, P. (2016). Predictive coding and representationalism. *Synthese*, 193, 559-582.
- Hobson, J. A. & Friston, K. J. (2012). Waking and dreaming consciousness: Neurobiological and functional considerations. *Progress in Neurobiology*, 98, 82-98.
- Hohwy, J. (2013). *The predictive mind*. Oxford: Oxford University Press.
- (2014). *The self-evidencing brain*. Noûs. <https://dx.doi.org/10.1111/nous.12062>.
- (forthcoming). The predictive processing theory and 4e cognition. In A. B. Newen & S. Gallagher (Eds.) *The Oxford handbook of cognition: Embodied, embedded, enactive and extended*. Oxford University Press.
- Hohwy, J., Roepstorff, A. & Friston, K. J. (2008). Predictive coding explains binocular rivalry: An epistemological review. *Cognition*, 108, 687-701.

- Kant, I. (1781/1996). *Critique of pure reason, abridged*. Indianapolis/Cambridge: Hackett Publishing Company.
- Kemp, C. & Tenenbaum, J. B. (2008). *The discovery of structural form* (pp. 10687-10692). 105.
- Kemp, C., Perfors, A. & Tenenbaum, J. B. (2007). Learning overhypotheses with hierarchical Bayesian models. *Developmental Science*, 10, 307-321.
- Kiefer, A. (2017). Literal perceptual inference. In T. Metzinger & W. Wiese (Eds.) *Philosophy of predictive processing*. Frankfurt am Main: MIND-Group.
- Langton, R. (1998). *Kantian humility. Our ignorance of things in themselves*. Oxford: Oxford University Press.
- Lupyan, G. (2015). Cognitive penetrability of perception in the age of prediction: Predictive systems are penetrable systems. *Review of Philosophy and Psychology*, 6, 547-569.
- Lyons, J. (2008). Evidence, experience, and externalism. *Australasian Journal of Philosophy*, 86, 461-479.
- Orlandi, N. (forthcoming). *Bayesian perception is ecological perception*. Philosophical Topics.
- Peirce, C. S. (1878/2011). How to make our ideas clear. In R. B. Talisse & S. F. Aikin (Eds.) *The pragmatism reader: From peirce through the present* (pp. 50-65). Princeton, NJ: Princeton University Press.
- Quine, W. V. O. (1969). Epistemology naturalized. *Ontological relativity and other essays* (pp. 69-90). New York: Columbia University Press.
- Ramsey, W. (2007). *Representation reconsidered*. Cambridge, MA: The MIT Press.
- Rescorla, M. (2015). Bayesian perceptual psychology. In M. Matthen (Ed.) *The oxford handbook of philosophy of perception* (pp. 694-717). Oxford: Oxford University Press.
- (2016). Bayesian sensorimotor psychology. *Mind & Language*, 31, 3-36.
- Sellars, W. (1956). Empiricism and the philosophy of mind. In H. Feigl & N. Scriven (Eds.) *Minnesota studies in the philosophy of science* (pp. 253-329). Minneapolis, MN: University of Minnesota Press.
- Seth, A. (2015). Inference to the best prediction. A reply to Wanja Wiese. In T. Metzinger & J. M. Windt (Eds.) *Open MIND* (pp. 4-20). Frankfurt am Main: MIND Group. <http://open-mind.net/papers/inference-to-the-best-prediction>.
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L. & Goodman, N. D. (2011). How to grow a mind: Statistics, structure and abstraction. *Science*, 331, 1279-1285.