Abstract

This chapter explores the consequences that the predictive processing framework has for the debate concerning what can enter into the content of perceptual experience. Within this debate, there are those who, on the one hand, are ‘conservatives’, claiming that only ‘low-level’ properties enter into the content of perceptual experience, while there are those who, on the other hand, are more ‘liberal’ and allow ‘higher-level’ properties. One recently expressed view is that the predictive processing framework dispenses with, or dissolves, the debate altogether. This chapter argues against this in favour of the view that predictive processing in fact supports liberalism. It defends this position against four illustrative objections.

Keywords

perception; content of perceptual experience; admissible contents; predictive processing framework; perceptual inference; content liberalism

9

What Can Predictive Processing Tell Us about the Content of Perceptual Experience?

Sam Wilkinson

1. Introduction

A central debate in philosophy of perception concerns the content of perceptual experience, which is, roughly, what is conveyed to the subject by, and only by, her perceptual experience. There are those who are 'conservatives' about this, claiming that only 'low-level' properties enter into this content (e.g. Tye 1995; Dretske 1995; Byrne 2016), and there are those who are more 'liberal', and allow 'higher-level' properties (e.g. Siewert 1998; Siegel 2006; Bayne 2009; Nanay 2011). In this chapter, I present the predictive processing framework (PPF) and explore the consequences that it may have for this debate.

One view about this (e.g. Lupyan and Clark 2015) is that the PPF dissolves this debate in the following way. The debate is about where a particular line (viz. between perception and cognition) should be drawn, but there is no such line according to the PPF. I argue against this and claim instead that the PPF is best understood as supporting liberalism. I end by defending this position against four objections.

2. The Content of Perceptual Experience (CPE) Debate

The debate we are interested in, which I am calling the 'Content of Perceptual Experience Debate', or CPE debate for short, concerns the types of properties that can enter into the

content of perceptual experience. What is the significance of saying that properties of a certain type enter into the content of perceptual experience? A concrete case (loosely borrowed from Nanay 2011) will be helpful here. Suppose you are looking at a green apple, and suppose that you know that it's a Granny Smith apple, and, furthermore, that it was grown in Chile. This apple has:

(i) a certain shape and colour.

(ii) the property of being an apple.

(iii) the property of being a Granny Smith apple.

(iv) the property of having been grown in Chile.

Neither side of the debate will want to say that all of those properties of the apple enter into the content of your perceptual experience. In particular, even the most liberal of liberals will likely accept that (iv) just isn't the right kind of property for your perceptual experience to convey. You may come to know that the apple has that property (viz. was grown in Chile), but you can't have known that *solely* on the basis of your perceptual experience. It can enter into the content of, say, a judgement or supposition, but not of a perceptual experience. And that is a large part of what this debate is supposed to be about: drawing a principled line between the sorts of properties that can enter into the contents of perceptual experiences on the one hand and judgements (or other less committal propositional attitudes) on the other.

What this amounts to, more practically speaking, is where you would attribute the 'blame' for any inaccuracy. Suppose, now, that you believe that the apple has all of (i–iv), and this time you are wrong about (iv): it wasn't grown in Chile, it was grown in France. It would be odd to say that, in such a situation, your perceptual experience alone had misled you. It was an error of judgement, either based on someone telling you the wrong thing, or you drawing a false inference (perhaps based on the time of year, you wrongly assumed a southern hemispheric provenance). Now, the conservative wants to say something similar

about (iii), and also about (ii). If you are wrong about it being a Granny Smith, or about it being an apple, but you are not wrong about it being a certain colour or shape (suppose it's a very realistic waxwork), this is not (not *ever*) a case of your perceptual experience alone leading you astray. Rather, it is the result of an incorrect inference based on an accurate perceptual experience. One reason why the conservative may think this is because your perceptual apparatus is working just fine and you're not subject to some kind of optical illusion. Moreover, the conservative may point to a more basic motivation for her position: the sorts of properties to which our eyes are sensitive are limited (and indeed we could even *define* 'low-level' properties as those properties). You can't have the property of being an apple enter into the content of a perceptual experience, because our eyes just aren't sensitive to apple-ness *per se*. To put this another way, our visual systems don't respond differentially to apples and to very realistic apple waxworks.

The liberal can counter this by pointing to an important distinction between being misled by your perceptual *experience* and errors in perceptual *processing*.[1] Just because there aren't errors in perceptual processing, it doesn't mean that your full-blooded perceptual experience isn't leading you astray. The liberal could appeal to the idea that even when you know it is a waxwork, if it is realistic enough, you have to override the experience of it as a real apple, an experience that wouldn't be had by a subject who had never encountered an apple before and had no conception of apples.

The classic arguments for liberalism exploit this idea that experiences had by those who have relevant knowledge or expertise differ phenomenologically from those had by those who lack it. In these 'contrast cases' you are invited to compare two cases where one would intuitively say that there is a phenomenological difference, in spite of the fact that the low-

[1] Arguably most optical illusions (e.g. the Muller-Lyer illusion) are cases of your visual system working fine, but your visual experience misleading you.

level properties represented (e.g. colour and shape) remain constant. It is then suggested that the phenomenological difference is best explained by the fact that in the two cases the 'high-level' properties represented are different because, in one of the cases, the high-level properties cannot be represented due to lack of knowledge, exposure or expertise. Contrast cases abound, but perhaps the classic example is from Siegel (2006):

> Suppose you have never seen a pine tree before, and are hired to cut down all the pine trees in a grove containing trees of many different sorts. Someone points out to you which trees are pine trees. Some weeks pass, and your disposition to distinguish the pine trees from the others improves. Eventually, you can spot the pine trees immediately: they become visually salient to you. Like the recognitional disposition you gain, the salience of the trees emerges gradually. Gaining this recognitional disposition is reflected in a phenomenological difference between the visual experiences had before and after the recognitional disposition was fully developed. (Siegel 2006, 491)

Siegel goes on to argue that this phenomenological difference is attributable to the fact that the kind property 'pine tree' enters, after recognitional expertise is gained, into the contents of the perceptual experience.

Responses to contrast cases can adopt two broad tactics in support of conservatism (as Siegel 2010 acknowledges):

1) Although the phenomenology has changed between the two cases, it is not perceptual phenomenology that has changed (it is, for example, emotional, affective or even cognitive phenomenology).
2) The phenomenology has changed, and it is indeed perceptual phenomenology, but that is because the low-level properties represented are different (for example, because it affects low-level processing, or how attention is directed). This amounts to

denying the founding assumption of the contrast: the low-level properties are not being kept constant after all.[2]

In short, conservatives usually acknowledge the phenomenological differences (although see Tye 1995) between the contrast case cases, but tend to give them an alternative explanation.

At this point it is worth mentioning that the desire to seek an alternative explanation that is consistent with conservatism must come from an independent motivation to adopt conservatism. One such motivation for conservatism, which I have already mentioned and which arguably makes it the default position, is that our eyes, ears, etc. seem to only respond to certain low-level features: our eyes respond to light, our ears to sound waves, etc. The framework that I am about to present goes some way towards undermining this motivation.

3. The Predictive Processing Framework

The *predictive processing framework* (PPF) makes us rethink both what the brain does in general, and the nature of perceptual experience in particular. Contrary to standard accounts, the brain does not take inputs from the outside world, process them, and pass them on (such a framework is exemplified by the use of 'box-and-arrow' diagrams in cognitive psychology). Instead, the brain is to be viewed as a prediction machine. Whenever information impacts on your sensory surfaces, it is already, even at the earliest stages, greeted by a prediction on the part of your nervous system. Your perceptual encounters with the world never occur in a vacuum, free of temporal context. Furthermore, what determines your conscious percept at a given time is not the inputs that causally and temporally precede the experience, but rather

[2] O'Callaghan (2011) uses this tactic for the auditory experience of speech, in particular to argue against the view that we can hear meanings.

the hypotheses that your brain has adopted in order to best predict future inputs. In other words, your brain is always staying one step ahead of the sensory manifold, and your experience is determined by your brain's best hypotheses. This framework has received empirical and theoretical support from a wide variety of sources and disciplines, and has far-reaching implications for how we are to think about all aspects of cognition (see Clark 2013, 2016; Hohwy 2013), including atypical cognition (for a treatment of psychosis, see Fletcher and Frith 2009; for a treatment of auditory verbal hallucinations, see Wilkinson 2014; for a treatment of autism, see Pellicano and Burr 2012).

Perhaps the simplest way to present the thinking behind the PPF is to reflect on two things: ambiguity and efficiency. The brain's main task is to figure out what is 'out there' based on what is impacting on the peripheries of the nervous system, and it has to do so as efficiently (namely, as quickly, accurately, and with as little energy expenditure) as possible. Since the incoming information is ambiguous (namely, compatible with different hypotheses) the brain has to use more than just the 'fit' to select one hypothesis over another. In particular, it has to use its 'priors', namely, its expectations (usually based on past statistical regularities). Thus with every stimulus, your brain is already greeting it with an expectation that helps to disambiguate it.

This tactic is not only a good way of disambiguating a stimulus, and of doing so quickly: it is also highly energy efficient. To see why, consider an analogy with data compression. Data compression is about minimizing 'information load', namely, that which is explicitly represented in a signal. A standard method for doing this between a given sender/receiver pair, is for the sender to only pass in the signal what is newsworthy, namely, what the receiver hasn't already predicted and hence can construct for itself. The same is said to occur between different regions of the nervous system: all that gets passed up is what the relevant part (the 'receiver') hasn't already predicted, which in the PPF is called 'prediction error'. In

turn, your brain is constantly trying to improve its predictions. Thus you get the dictum: *all the brain ever does is minimize prediction error*.

The extent to which this very strong and universal claim is true is something that I will put to one side. However, there is evidence converging on the idea that the brain is in the business of doing something like this (see Clark 2013 for a presentation of some of this evidence from a wide variety of different disciplines), and what is important for my purposes is that this has far-reaching consequences for how we are to think about perceptual experience. What is of vital importance here is that things don't impact on our nervous systems out of context, against a static background: our nervous systems are in a constant fluctuating state of anticipation. Your experience is not constructed out of sensory inputs, but is rather the result of a dynamic, predictive hypothesis-building process.

One important aspect of the PPF is that the predictive hypotheses that the brain selects are hierarchically organized, with the hypotheses of one level providing the data for the next. 'Higher' parts of the hierarchy are, roughly, those parts that are further away from the sensory surfaces. These tend to be at lower temporal frequencies (longer timescales), and at higher levels of abstraction. 'Lower' parts of the hierarchy are closer to the sensory apparatus. These tend to be at higher temporal frequencies (shorter timescales), and at low levels of abstraction. They correspond, for example, to early stages of visual processing: your brain's early statistically driven attempts to make sense of noisy inputs. Of course, in order to express these neurally encoded predictions we need to use descriptions in natural language (in this case English), but there is nothing linguistic about the priors/hypotheses ('Light tends to come from above'/'This is a face') themselves (indeed, it is a matter of some dispute whether they themselves are representational at all (see Gladziejewski 2016 for arguments for, and Orlandi 2016 for arguments against)).

4. Counterfactual Predictions and 'Seeing As'

According to the PPF, perceptual experience consists in your brain selecting the hypotheses that generate the best predictions of imminent sensory data. However, the perceptual conditions at any given time hugely underdetermine the hypotheses that could be appropriately selected. The nature and complexity of the hypotheses is left open, as is, in a related manner, the nature and reach of the predictions that such hypotheses may generate. If you put an apple in front of a fully sighted human being in good lighting conditions, then the hypotheses that their brain selects could (over and above the low-level hypotheses that settle questions about colour and shape) vary enormously. Perhaps the person has never seen an apple before, so they experience it as 'That thing, which has that shape, appears to have solidity etc. . . .' Perhaps the person has never seen an apple, but recognizes it as a fruit ('That strange fruit I've never seen before'). There is no limit to the type and degree of variation we could dream up. For example, other variation may not come from lack of exposure or expertise (which would instil Siegel's 'recognitional disposition') but from background information about that very specific perceptual encounter. Perhaps the person has good reason to think (e.g. were told by a reliable source) that, in spite of appearances, it's a waxwork apple. Yet other, relatively short-lived variations may arise from the person's practical concerns. Perhaps the person is an impressionist painter and momentarily brackets the apple-ness simply to focus on the arrangement of colour in her visual field.

In mainstream analytic philosophy of perception these variations would be chalked up to differences in 'seeing as', which amounts to the application of different concepts to, or different conceptualizations of, the same sensory data. That is not strictly incompatible with the PPF. Perhaps, for example, one can think of 'concepts' as the very same thing as the 'hypotheses' that the PPF posits. However, two caveats should be flagged. First, according to the PPF, such 'concepts' aren't to be seen as applied to perceptual *experience* but rather as

generating (or determining) the experience.[3] Second, the PPF will inevitably explain the differences in concepts or conceptualizations in terms of different predictions. This could play out in the following way.

Everything has to boil down to predictions, and the reach of these predictions can—and does—vary enormously. To borrow an insight from Seth (2014), our hypotheses are *counterfactually rich*, and can vary in their counterfactual richness. What he means by this is that my brain doesn't simply have predictions about what is going to happen, namely, how the sensory manifold *will* change given what is likely to happen in the near future. There is also (at least in the human brain) a wealth of prediction about how the sensory manifold *would* change given relevant counterfactual circumstances—circumstances that needn't happen and may even be unlikely to happen. Another way of putting this is that my nervous system isn't simply content with superficially predicting what will happen next, but seeks to comprehend more fully the statistical structure of the world, which involves, in part, the tacit positing of the underlying natures and dispositions of things, even if those natures and dispositions are never explicitly revealed, since their being revealed might involve an excessively complicated and/or unlikely set of counterfactual circumstances. This deeper and more costly enterprise clearly amounts to a wise long-term investment, since it prepares you for a causally complex and hard-to-predict world.[4]

My experience of an apple as an apple is partly down to my brain's predictions about what *would* happen *if* I were to do any number of things (walk around it, pick it up, eat it etc.) or if any number of things were to be done to it (e.g. if it were smashed with a

[3] This may be in keeping with the views of 'conceptualists' about perceptual experience (e.g. Brewer 1999; McDowell 1994).

[4] Indeed, nowhere is causal complexity greater, and predictability more important, than in the social realm (as it involves the perception of other agents).

sledgehammer), even if those possibilities aren't likely to be actualized (e.g. I have every intention of staying still, of continuing to view the apple from this angle, and there is no threat of the apple being smashed with a sledgehammer). My nervous system's 'appreciation' of these unactualized possibilities structures my experience and enables it to delve below the surfaces of objects to the postulated nature of those objects. Notice that the apple could be in a glass cabinet, such that I never get to confirm or falsify predictions that determine my experience of it as a real apple. Suppose that, unbeknownst to me, the object in the cabinet is a highly realistic waxwork apple. What it is for me to see it, wrongly in this case, as a real apple is the plethora of counterfactual predictions pertaining to real apples (e.g. that I can eat it; that it would smash in a certain way etc.).

Some readers may have noticed a similarity between these counterfactually rich predictions and so-called sensorimotor contingencies (O'Regan and Noe 2001). Indeed, this is a similarity that Seth (2014) notes, and uses to great effect (as Noe himself does (2005) within an enactivist framework) in making sense of perceptual presence (and variations therein).[5] However, it is crucial to see that not all counterfactual predictions are sensorimotor contingencies. They are not necessarily latent expectations about what would happen if *I* were to *do* something. Many of the relevant expectations are to do with what will happen to the object in certain circumstances that do not involve my acting at all. This is not to contradict Seth (2014) since I think it is likely that the predictions that are relevant specifically to perceptual presence are indeed action-centred counterfactual predictions (viz. sensorimotor contingencies) (see Wilkinson 2020). There is, however, rather more to perceptual experience than perceptual presence.

[5] 'Perceptual presence' refers to the experience of something as being present to you, as part of your world.

5. Does the PPF Dissolve the CPE Debate?

What does the PPF have to contribute to the CPE debate? One view, expressed by Lupyan and Clark (2015), but also easily derivable from what other PPF theorists say (e.g. Fletcher and Frith 2009; Hohwy 2013) is that the PPF dissolves the CPE debate altogether.[6] The debate is based on a misguided picture of the human mind. The debate is about where you draw the line between the perceptual and the cognitive, and what the PPF tells us is that there is no such line to be drawn. What you have all the way 'up', and indeed all the way 'down', is a delicate dance between perception and cognition.

Such theorists will want to deny, for example, that there is a principled distinction to be drawn between perception and belief on the grounds that there is no *fundamental* difference between my visual system's 'beliefs' about what is out there, and *my* belief that my train leaves at 10:30. It is simply a difference to do with where in the hierarchy these are located: one is 'high up', abstract, and operates over longer timescales; the other is 'low down', concrete, and operates over shorter timescales. But these differences simply fall out of the fundamental nature of the hierarchy: they are, at bottom, the same kind of phenomena.

If we apply this line of thinking to Siegel's contrast cases, and to liberalism more generally, the question of whether, e.g., natural kinds enter in the content of perceptual experience is to miss the point. My nervous system has a certain hierarchical predictive structure that reflects my understanding of the world, and it both shapes and is shaped by my engagement with the world. Is the phenomenology of my experience of a Granny Smith apple, as an apple expert, different from the experience of a naïve observer because of something perceptual, or because of a judgement made on the basis of a percept? For the PP

[6] Fletcher and Frith (2009) in particular argue that hallucinations and delusions in psychosis are fundamentally the same phenomenon, simply operating at different levels in the hierarchy.

theorist, it is based on something judgement-like in my nervous system. But this is no victory for the conservative, because, then again, the same applies to my appreciation of the object's colour and shape. It is judgement all the way down.

I want to argue against this view. In spite of similarities (e.g. they have mind-to-world direction of fit) there is a major functional difference between these two kinds of 'judgements/beliefs', namely, those that I make/have and those that my nervous system makes/has. That difference is to do with the way in which they are formed or updated. I just need to hear an announcement that my train has been delayed by ten minutes in order for me to update my belief about my train's departure. To get my visual system to update its 'belief', e.g., that those two lines in the Muller-Lyer illusion are in fact the same length, is rather more difficult. It would, if it turned out to be possible at all, require a lengthy process of gradual perceptual relearning.[7] And isn't that where we should naturally draw the line? Belief is about what I believe as a rational agent, and, as a result, can be responsive to one-off pieces of evidence, whereas perception is about what a part of my nervous system 'believes' as a result of statistically driven priors, and it may continue to 'believe' this whatever I myself come to discover (for example, by measuring the lines of the Muller-Lyer illusion).

6. The PPF as Supporting Liberalism

I now want to take things one step further. Not only does the PPF not dissolve the CPE debate, it can also be interpreted as offering support to liberalism. The most important lesson from the PPF is that our nervous system constructs hypotheses that delve beneath the surfaces of what is explicitly, sensorially presented, in order to best predict what will or might (or

[7] Indeed it is believed that those who have developed in non-rectilinear environments are less susceptible (or not at all) to the Muller-Lyer illusion (Berry 1968).

could) be sensorially presented in the future. These hypotheses will be tweaked and updated if they fail to do this. However, the crucial move for my purposes lies in realizing that *it is the hypotheses themselves that ultimately determine the content of experience, and not the sensory inputs that primarily serve to keep them accurate and anchored to the world*.[8] In light of this theoretical perspective, the kinds of properties that can enter into the contents of experience are not restricted to the sorts of properties that my sensory apparatus can detect. To think that such a restriction applies would be to conflate the sort of data that can falsify or confirm a hypothesis, with the content of the hypothesis itself. Typically, a hypothesis is rich and theory-laden, whereas the data against which the predictions it generates are compared are in and of themselves sparse, ambiguous, and open to innumerable interpretations. However, as we are about to see, this is not to say that there are no restrictions on the kinds of things that can enter into the contents of perceptual hypotheses (as opposed to more 'cognitive' hypotheses).

Indeed, the reason why we experience a rich, structured and meaningful world is precisely because the hypothesizing of our nervous systems far outstrips, delves far beyond,

[8] Of course, the importance of this 'anchoring' shouldn't be downplayed. (After all, when it goes wrong in different ways, you get different kinds of disconnection from reality.) In a case of perception, the experience is causally dependent on signals from the outside world (you wouldn't get *that* experience without them) but it is not (psychologically, proximally) determined by them (or, perhaps better, constituted by them), in that in principle (though not in practice) you *could* get that experience without them. In any case, we are not interested here in perceptions vs. hallucinations; we are interested in the richness of perceptual experiences, and that richness comes from the hypotheses selected. We can imagine super-advanced human beings who, with the same sensory apparatus, when engaged in the very same perceptual encounter with the world, have an experience of unfathomable richness compared to our own (or, conversely, others whose experience is impoverished).

the sensory data at their disposal. Indeed, one might even go as far as to say that, not only does the PPF remove the restriction of sensory detectability, it also makes nonsense of the very idea of raw sensation at the experiential level. To the extent that any sensation is consciously experienced, it has already been accounted for *to some extent* by a top-down hypothesis. The closest approximation in experience to the notion of raw sensation is sensory data that has only been accounted for by a relatively low-level hypothesis.

In light of all of this, one might ask: What kinds of properties *aren't* allowed in experience? As we've seen, the relevant restrictions are no longer about the kinds of things that your sensory surfaces are sensitive to. But this is not to say that there are no restrictions at all. These restrictions concern the sorts of hypotheses that your nervous system can generate at perceptual timescales. Two crucial points are embedded in this: (i) the hypotheses that *your nervous system* can generate, and (ii) at *perceptual timescales*.

To clarify (i), it helps to rehearse a line of criticism. Are we not, the critic might ask, saying that these are the contents of *judgements* (albeit perceptual ones)? The reason why this is not the case is because it is crucial to distinguish judgements that *I* make, from the 'judgements' that my *nervous system* (or a part of my nervous system) makes.[9] The latter 'judgements' are what determine the content of my experience, and, crucially, can be at odds with the judgements that I make. Nowhere is this more clearly seen than in cases where we knowingly experience illusions. When, for example, I experience the Muller-Lyer illusion, my nervous system 'judges' that the lines are different lengths, whereas I judge them to be the same length (for example, because I have measured them). What we typically think of as a perceptual judgement is when I endorse the 'judgement' that my nervous system has made, namely, the lower-level 'inference' that generates the perceptual content.

[9] Note the deliberate use of so-called scare quotes where I'm using personal-level vocabulary to capture something subpersonal but functionally similar.

Regarding (ii), hypotheses generated at 'perceptual timescales' are hypotheses that are relevant to my online engagement with a given stimulus. Thus there are all sorts of judgements, ways that we take the world to be, that are too abstract, or general, or temporally distant or distributed to be perceptual hypotheses. For example, my belief that Paris is the capital of France, or that my train leaves at noon tomorrow, or, indeed, that that apple was grown in Chile, *these* hypotheses (and they are in an important sense hypotheses) are not directly relevant to my online engagement with the object in question. Historical and political properties (like having been grown somewhere, or being the capital of somewhere) can certainly be of relevance, in the right contexts, to my engagement with something, but they are not the sorts of things that make a difference to my there-and-then, stimulus-dependent engagement with the object perceived. In contrast, the hypothesis that the thing in question is a real apple is very relevant to my there-and-then engagement with it.

What about the property of being a Granny Smith? Such a hypothesis could be at a perceptual timescale (in fact, there's no reason to think that it wouldn't be at the same timescale as the real apple hypothesis). However, what is less clear is whether the Granny Smith content can be part of my nervous system's hypothesizing, rather than my own. It seems that an ability to recognize Granny Smiths is not sufficient for this. This could still be the content of an inferential judgement (as in the early stages of Siegel's example). But I see no reason why a high degree of expertise couldn't enable the Granny Smith hypothesis to be something that my nervous system automatically selects in appropriate circumstances. The property of having been grown in Chile, however, is inappropriate in a number of ways: it is too abstract, too distant in both space and time, and so on.

7. Four Illustrative Objections

I'd like to present four illustrative objections to what I have argued. Each of them can be addressed in ways that usefully clarify the position I'm presenting. Since my claim is reliant on the viability of the PPF, the objections don't question that. They rather cast doubt on my suggestion that the PPF supports liberalism. The first two will be familiar, since they are effectively the same as the two standard responses to contrast cases. The last two are new.

7.1 Phenomenological Effects Are Attributable to Top-down Effects on the Experience of Low-Level Properties

When my brain adopts one multilevel hypothesis instead of another competing one that would be equally acceptable in light of current perceptual conditions (e.g. the real apple vs. the waxwork hypotheses) what we get is indeed a change in phenomenology. However, this change in phenomenology is not because high-level properties enter into the contents of perceptual experiences but rather because, as the PPF so strongly emphasizes, there are pervasive top-down influences on how the low-level properties are experienced. In other words, top-down effects change the properly *perceptual* content, but simply by changing which low-level properties are experienced. That, in itself, the objection goes, is perfectly consistent with conservatism, given that it is only these now-changed low-level properties that enter into perceptual experience.

That is true. However, this objection underestimates both the nature and the reach of top-down influences within the PPF. It is the top-down predictions (the hypotheses) that determine your experience, not the bottom-up sensory prediction error, even at the front-line sensory surfaces. In other words, even the 'lowest level' of perceptual experience is only possible thanks to top-down processing. As we've seen, within the PPF there is no such thing as 'raw sensory experience'. To the extent that there is any sensory data at all, it is already, even at the earliest stages, being greeted by expectations on the part of your nervous system,

and the 'rawest' sensory *experience*, if it is to be considered an experience at all, is simply the result of a low-level hypothesis. So why (and indeed where) should you draw the line between the low-level and high-level aspects of experiences if both are ultimately generated top-down? Short of answering this, you should either say that neither the lower-level nor the higher-level aspects can feature in perceptual content, in which case nothing can, or you allow that both can. In short, if the PPF is correct, either you do away with the very notion of perceptual content or you should be liberal about it. As I have already argued, in the context of discussing illusions that remain in the absence of judgements to the contrary, I think there are good reasons to retain a notion of contents that are perceptual as opposed to more robustly cognitive (i.e. associated with judgement).

7.2 Phenomenological Effects Are Attributable to Non-perceptual Phenomenology

The second objection is that, when our experience is of these higher-level things in virtue of its predictive structure that might be part of the overall experience, but not of the *perceptual* experience. Two obvious candidates for such non-perceptual phenomenological contributions to the experience are, first, affective phenomenology and cognitive phenomenology. I examine these in turn.

So, can the relevant phenomenological contribution be seen as a matter of accompanying affect rather than perceptual content *per se*? An initial response to this is that, indeed, perceptual experiences, and especially anticipatory aspects of those experiences, are often affectively charged, but that doesn't make them any less perceptual. Not only are even the most mundane forms of perceptual experience subtly affective, affect is also constantly modulating perceptual processing (e.g. Phelps, Ling and Carrasco 2006; Villeumier and Driver 2007). Affect doesn't merely accompany perceptual processing. This phrasing would

suggest that you could remove it and leave the perceptual processing intact. In truth, affect and perception are profoundly intertwined. This is evidenced, for example, by rare cases of brain damage where affective disruptions correspond to profound changes in what is experienced. These changes can lead the subject to claim that loved ones have been replaced by identical-looking impostors, as in the Capgras delusion, or even, as in the case of the Cotard delusion, that one is dead and that the world isn't real.

What about cognitive phenomenology? Here it depends on what you mean by this. Sometimes 'cognitive phenomenology' is taken to refer to the phenomenology of 'cognitive' (as opposed to 'conative') propositional attitude states and events. Thus we would be talking here about a phenomenology of belief or judgement. This doesn't seem like a good candidate. Not only is the very existence of cognitive phenomenology in this sense highly contentious (see Bayne and Montague 2011), but the relevant aspects of perceptual content are supposed to be independent of what I judge. I can continue to experience something as edible or fragile or a real apple, to the extent that my nervous system continues to select that hypothesis, even though I myself judge this to be inaccurate.

Another option is to think of 'cognitive phenomenology' as meaning the 'phenomenology of cognition' where cognition is broadly construed so as to capture the subject matter of the cognitive sciences. This would not be about beliefs and judgements, but may include, for example, conceptualization and categorization. In this case, the answer is that, yes, the phenomenological difference between seeing something as a real and a fake apple is a matter of cognition, but cognition in this sense would in any case be an integral part of perception, and the generation of perceptual content.

7.3 The PPF Is Entirely Consistent with Conservatism

My argument rests on the idea that the PPF supports liberalism because, according to the PPF, it is not the sensory input that determines perceptual content, but the hypotheses used to adequately predict the dynamics of future input. However, the objection goes, nothing prevents someone from adhering to the PPF whilst claiming that it is only *perceptual* hypotheses that determine perceptual content and *these* concern only low-level properties. In other words, only a liberal version of the PPF (rather obviously) supports liberalism, whereas there is an equally coherent version that supports conservatism. In short, the PPF does not in and of itself support liberalism. This objection is extremely helpful in allowing me to clarify the precise nature of my central claim, on which I'd like to say two things.

First, I don't want to say that the PPF is *inconsistent* with conservatism. One could imagine a hierarchically arranged predictive processing architecture of a conservative sort. Indeed, perhaps some actual organisms have nervous systems that approximate this. My claim is simply that human beings aren't like this. My reasons for making this claim are partly due to the kinds of organisms that we are: we are, for example, deeply social and linguistic animals that develop rapid-timescale and inflexible expertise (as reflected within the PPF by the sorts of hypotheses that our nervous systems generate) about domains of reality that are underdetermined by, and far outstrip, the merely sensory. In a related vein, the claim that we have a perceptually conservative nervous system strikes me as phenomenologically implausible. More mundane liberal intuitions are onto something: we don't just see colours and shapes, but apples, chairs, faces, etc.

The second thing I'd like to say is that, while not strictly *incompatible* with conservatism, in focusing on hypothesis generation rather than sensory input for the determination of perceptual content, the PPF undermines some of the motivation for conservatism (namely that our sensory surfaces are restricted in the kinds of things that they

can respond to). And with this motivation undermined, why would you fight the various arguments for liberalism?

What is more, conservatism is further undermined when you reflect on whether it is plausible that the human nervous system would restrict itself to generating predictive hypotheses about the here and now, at rapid, action-guiding timescales, that are only about (what gets called in the CPE debate) low-level properties. Clearly delving beyond these superficial sensory features is an integral part of how we experience the world and is a vital part of our predictive success, success which concerns what is right here in front of me, to be responded to now. I see no reason why this online, here-and-now informational state should not be called a perceptual state.

7.4 What Modality Are These So-Called Perceptual Experiences Happening In?

Perceptual experiences happen in a given modality. I see colours, hear sounds, etc. But when I experience edibility, or fragility is this something that I specifically see? And how is it that my eyes are supposed to respond differentially to these properties? Something that's edible or fragile can impact upon my visual system in exactly the same way as something that is not those things.

A potential response to this, again, involves a clarification of the relationship between sensation and perception. Sensation is necessary for perception, but certainly not sufficient.[10] I become perceptually informed about the world in virtue of my sensory apparatus. But much

[10] Although this necessity is probably only physical rather than metaphysical.

of this is cross-modal and multimodal.[11] The natural way to talk about these cases is not to say that I literally see something's edibility, but rather that I experience that something is edible in virtue of seeing it. But it could be in virtue of touching it, or smelling it, or through a combination of senses. Perhaps, in light of the PPF, one should not *really* talk about perception of high-level properties in a different way to how we talk about perception of low-level properties. That is to say, perhaps we don't literally see shapes, or even colours, but rather perceive that things have certain shapes or colours in virtue of sight. We don't hear sounds, but rather perceive that certain sounds were produced (or indeed that there was a distant car crash) in virtue of audition. Remember, after all, that it is the perceptual hypotheses that your nervous system selects in order to best predict sensory input that determines the content of your perceptual experience, and not the sensory input itself.

8. Conclusion

I explored the consequences that predictive processing has for how we are to think about the contents of perceptual experience. Those who are sympathetic to the PPF have tended to think that the debate between liberals and conservatives goes away. In contrast, I have argued that the PPF should be understood as supporting liberalism. This support comes in two stages.

First, motivations for conservatism that come from constraints about the kinds of properties that our sensory organs are sensitive to are undermined since the emphasis in the PPF is on the hypotheses that best predict sensory input, rather than the input themselves

[11] And in any case, there is very little consensus about how to individuate modalities in the first place (see Macpherson 2011).

(indeed, the idea of a 'raw sensory input' makes no sense within the PPF, at least not as a component of experience).

Second, the challenge then becomes how one should draw the line between the hypotheses that count as perceptual and those that are more cognitive (belief-like). At this point I have appealed to two things. First, the fact that my beliefs and the 'beliefs' of my nervous system are rather different, and the way they function (e.g. in formation and update) are very different. It is only the latter that should be thought of as perceptual.[12] The second thing is that perceptual hypotheses are about what is around me in the here and now, and serve to guide potential actions at relatively short timescales.

Funding

European Research Council (XSPECT—DLV-692739)

References

Bayne, T. 2009. 'Perception and the reach of phenomenal content'. *Philosophical Quarterly* 59 (236): 385–404.

Bayne, T. and Montague, M. (eds.). 2011. *Cognitive Phenomenology*. Oxford and New York: Oxford University Press.

Berry, J. W. 1968. 'Ecology, perceptual development and the Müller-Lyer illusion'. *British Journal of Psychology* 59 (3): 205–10.

[12] Or at least a subset should be. There are hypotheses that my nervous system adopts to glean information not directly about the word, but to flag bodily change, some of which might be indirectly informative about the world. This is where the PPF accounts for emotion (Seth 2013). This may not count as perception, but there are clearly aspects of emotional experience that are very perception-like.

Brewer, B. 1999. *Perception and Reason*. Oxford: Oxford University Press.

Byrne, A. 2016. 'The epistemic significance of experience'. *Philosophical Studies* 173 (4): 947–67.

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

Clark, A. 2016. *Surfing Uncertainty*. Oxford: Oxford University Press.

Dretske, F. 1995. *Naturalizing the Mind*, Cambridge, MA: MIT Press.

Fletcher, P. C. and Frith, C. 2009. 'Perceiving is believing: A Bayesian approach to explaining the positive symptoms of schizophrenia'. *Nature Reviews Neuroscience* 10(1): 48–58.

Gładziejewski, P. 2016. 'Predictive coding and representationalism'. *Synthese* 193: 559–82.

Hohwy, J. 2013. *The Predictive Mind*. Oxford: Oxford University Press.

Lupyan, G. and Clark, A. 2015. 'Words and the world: predictive coding and the language-perception-cognition interface'. *Current Directions in Psychological Science* 24: 279–84.

Macpherson, F. 2011. 'Individuating the Senses'. In Fiona Macpherson (ed.) *The Senses: Classic and Contemporary Philosophical Perspectives*. Oxford University Press.

McDowell, J. 1994. 'The Content of Perceptual Experience'. *Philosophical Quarterly* 44: 190–205.

Nanay, B. 2011. 'Do we see apples as edible?' *Pacific Philosophical Quarterly* 92 (3): 305–22.

Noë, A. 2005. 'Real Presence'. *Philosophical Topics* 33 (1): 235–64.

O'Callaghan, C. 2011. 'Against hearing meanings'. *Philosophical Quarterly* 61 (245): 783–807.

O'Regan, J. K. and Noë, A. 2001. 'A sensorimotor account of vision and visual consciousness'. *Behavioral and Brain Sciences* 24 (5): 883–917.

Orlandi, N. 2016. ‘Bayesian perception is ecological perception’. *Philosophical Topics* 44 (2): 327–51.

Pellicano, E. and Burr, D. 2012. ‘When the world becomes “too real”: a Bayesian explanation of autistic perception’. *Trends in Cognitive Sciences* 16 (10): 504–10.

Phelps, E.A., Ling, S., and Carrasco, M. 2006. ‘Emotion facilitates perception and potentiates the perceptual benefits of attention’. *Psychological Science* 17: 292–9.

Siegel, S. 2006. ‘Which Properties are Represented in Perception?’ In T. Gendler Szabo and J. Hawthorne (eds.) *Perceptual Experience*. Oxford: Oxford University Press, 481–503.

Siewert, C. 1998. *The Significance of Consciousness*. Princeton: Princeton University Press.

Seth, A. K. 2013. ‘Interoceptive inference, emotion, and the embodied self’. *Trends in Cognitive Sciences* 17 (11): 565–73.

Seth, A. 2014. ‘A predictive processing theory of sensorimotor contingencies: Explaining the puzzle of perceptual presence and its absence in synaesthesia’. *Cognitive Neuroscience* 5 (2): 97–118.

Tye, M. 1995. *Ten Problems of Consciousness*. Cambridge, MA: MIT Press.

Vuilleumier, P. and Driver, J. 2007. ‘Modulation of visual processing by attention and emotion: windows on causal interactions between human brain regions’. *Philosophical Transactions of the Royal Society B: Biological Sciences* 362: 837–55.

Wilkinson, S. 2014. ‘Accounting for the varieties and phenomenology of auditory verbal hallucinations within a predictive processing framework’. *Consciousness and Cognition* 30: 142–55.

Wilkinson, S. 2020. ‘Distinguishing Volumetric Content from Perceptual Presence within a Predictive Processing Framework’. *Phenomenology and the Cognitive Sciences* 19 (4):791-800.