

Three Problems for the Predictive Coding Theory of Attention

Abstract: While philosophers of science and epistemologists are well acquainted with Bayesian methods of belief updating, there is a new Bayesian revolution sweeping neuroscience and perceptual psychology. First proposed by Helmholtz, predictive coding is the view that the human brain is fundamentally a hypothesis generator. Though predictive coding has most prominently offered a theory of perception – the bulk of the empirical support for the theory also lies in this domain – the Bayesian framework also promises to deliver a comprehensive theory of attention that falls out of the perceptual theory without the need for positing additional machinery.

The predictive coding (PC) theory of attention proposed by Feldman & Friston (2010) and defended by Hohwy (2012, 2013) is that attention is “the process of optimizing precision of prediction errors in hierarchical perceptual inference” (Hohwy 2013, p.195). Prediction errors are measurements of the difference, or mismatch, between predicted and actual evidence. Expected precisions are a measure of how reliable, or precise, we expect the prediction error signal to be in a given context: how likely is it in a given situation that the incongruent data constitutes legitimate prediction error as opposed to noise? On this picture, attention has the functional role of guiding perceptual inference by directing processing resources towards the prediction errors with the higher expected precisions.

We argue here that this theory of attention faces significant challenges on three counts. First, while the theory may provide a successful account of endogenous spatial attention, it fails to model endogenous feature-based attention: for attention to be driven by expectations of precision, then it has to be driven to an area where a large prediction error is generated. However, this is the inverse of what is needed to drive attention towards the relevant object. We further consider whether Clark’s (2013) proposed ‘provisional detection’ solution to a similar problem raised by (Bowman, Filetti, Wyble, & Olivers, 2013) can be understood along the lines of ‘gist perception’ (Bar, 2003), and whether this resolves the issue.

Second, it is unclear how the theory may accommodate non-perceptual forms of attention such as attention to one’s thoughts. The PC theory of attention is committed to the claim that attention just is the amplification of gain on prediction error, and that this is driven by expected precision. So the proposal would be that we pay attention to our thoughts when we expect them to be precise. However, this proposal remains to be filled out. Do we expect our thoughts to be more precise on some occasions rather than others? If so, what learned causal regularity underlies this expectation?

Third, it fails to accommodate the influence of affectively salient objects or high cost

situations in guiding and capturing attention. This points to a more general need to integrate both agent-level preferences and the cost of false negatives and false positives into the model, such that standards for expected precision can be adjusted. The challenge for the PC theory of attention is then to accommodate these additional influences on attention in terms of expected precisions.

Predictive coding in neuroscience

While philosophers of science and epistemologists are well acquainted with Bayesian methods of belief updating, there is a new Bayesian revolution sweeping neuroscience and perceptual psychology. First proposed by Helmholtz (2005), and with formal roots in signal processing data compression strategies (Shi & Sun, 1999) and pattern recognition in machine learning (Bishop, 2006), predictive coding is the view that the human brain is fundamentally a hypothesis generator.

On this view, the processes by which the brain tests its self-generated hypotheses against sensory evidence are seen as conforming to a hierarchical Bayesian operation; each level of the hierarchy involves a hypothesis space, with higher levels generating hypotheses about more complex and slower regularities as compared to the lower levels. The higher-level hypothesis spaces serve to generate and constrain the lower-level hypothesis spaces, thus enabling the lower-levels to predict the evidence. When there is a mismatch between the predicted and actual evidence, a prediction error is produced and is relayed up the hierarchy, where it is used to revise the hypothesis. Through the iterative interaction between top-down signals (which encode predictions) and bottom-up signals (which encode prediction error) the generative models that can predict the evidence most accurately are selected.

Given the crucial role of sensory evidence in supervising the hypothesis testing process, it is no surprise that the view has garnered the most significant empirical support as a theory of perception (Hohwy, Roepstorff, & Friston, 2008; Huang & Rao, 2011; Stefanics, Kremlacek, & Czigler, 2014). Nonetheless, increasing numbers of neuroscientists are also adopting the predictive coding framework in some capacity in order to elucidate attention, action (Berniker & Kording, 2011; Friston, Daunizeau, Kilner, & Kiebel, 2010; Körding & Wolpert, 2006), dreaming (Hobson & Friston, 2012), schizophrenia (Adams, Perrinet, & Friston, 2012; Horga, Schatz, Abi-Dargham, & Peterson, 2014; Wilkinson, 2014), interoception and the emotions (Seth & Critchley, 2013; Seth, Suzuki, & Critchley, 2011).

The predictive coding theory of attention

The predictive coding (PC) theory of attention proposed by Feldman & Friston (2010)

and defended by Hohwy (2012, 2013) is that attention is “the process of optimizing precision of prediction errors in hierarchical perceptual inference” (Hohwy 2013, p.195).¹ Prediction errors are measurements of the difference, or mismatch, between predicted and actual evidence. Expected precisions are a measure of how reliable, or precise, we expect the prediction error signal to be in a given context: how likely is it in a given situation that the incongruent data constitutes legitimate prediction error as opposed to noise?

Optimizing expected precisions is the process of guiding hypothesis revision by directing processing resources towards the prediction errors with the higher expected precisions – we attend to what is expected to be the most informative, and this information is used to preferentially revise our perceptual hypotheses. Such a practice allows us to avoid the potentially disastrous consequences of revising our hypotheses on the basis of noise-induced prediction error.

On this picture attention has the functional role of guiding perceptual inference by directing processing resources towards the prediction errors with the higher expected precisions. Again, this results in the minimization of prediction error, though attention is concerned only with expected precision of prediction error and not directly with the accuracy of the hypotheses. However, because the estimation of expected precisions is a fundamental aspect of perceptual inference, then so is attention. While the account is meant to be a comprehensive theory that encompasses both endogenous and exogenous attention, this paper will focus primarily on the former.

In exogenous (bottom-up) attention, the presentation of a contextually salient stimulus results in an abrupt and large prediction error. This large prediction error will draw one’s attention to the unattended stimulus because of a learned causal regularity that stronger signals are more precise (Hohwy 2013). Given that on the PC theory signal is defined as prediction error, then a strong signal will be one that has a large prediction error. Large prediction errors count as stronger signals because they are expected to be more informative. Given that larger signals are expected to be more precise for this reason, the gain or amplitude of this large prediction error will be enhanced (which just amounts to paying attention to the stimulus). Attention will then cause the hypothesis to be revised preferentially in light of this prediction error.

Endogenous (top-down) attention can be understood in terms of a conscious decision to attend to a given object or spatial region, or it can be understood more minimally in terms of endogenous cueing that requires agent interpretation. Using the classic Posner paradigm as an illustrative device, the PC theory of attention provides the following account of endogenous spatial attention. First, through repeated trials the subject learns that when an arrow is shown pointing to a given area on a computer screen, an object will likely appear

¹ Other theorists with similar views include (Rao & Ballard, 2004; Spratling, 2008; Summerfield & Egner, 2009).

in that area. This learned causal regularity is a contextually mediated expectation for precision: when there is an arrow pointing towards a given location, the prediction error that will subsequently be produced by the appearance of the object in that location is expected to be precise, or reliable. Second, suppose an arrow appears on the screen, pointing to the bottom right corner. This causes two things to happen: (i) the prior probability of there being an object in the bottom right corner goes up; and (ii) the gain from the prediction error issuing from this region is increased (this is tantamount to saying that one pays attention to the bottom right corner, as gain is identified with attention). Third, when the stimulus appears it is perceived more rapidly, for two reasons: the gain on the prediction error for this spatial region is enhanced, making it such that this prediction error drives the revision of the hypothesis; and the higher prior probability accorded to the hypothesis that an object will appear in this corner makes it the case that this hypothesis is more likely to be amongst those selected to drive perceptual inference in the first place. It allows the perceptual inference process to begin with a more accurate hypothesis, and so spend less time in revisions.

A problem for the PC theory of endogenous attention

Note that the arrow cue does not predict which object will appear (unless a hypothesis has been formed for this as well via conditioning, such as that dots are likely to appear after arrow cues). However, Hohwy claims that the same sort of explanation can be applied to feature-based endogenous attention (2013, p.196). In such cases, an object ‘pops out’ of a scene when one has been given the task of looking for it. How might this work? It is crucial that it do so, as many cases of endogenous attention are those involving searches for certain features or objects over others. However, it is unclear how the account is supposed to go, given that attention must be driven by high expected precision.

To illustrate the problem that arises for the PC account, take the case of searching for one’s keys. What are the relevant precision expectations driving attention? They cannot be spatial – one doesn’t have high expected precision for any particular spatial region (beyond a few general expectations, such as that one’s lost keys typically won’t be found hanging from the ceiling). Perhaps one begins with a high expected precision for any prediction error generated by the hypothesis ‘this item is my keyset’. Certain features – silver, key-shaped, jangling if moved, will drive lower level hypotheses, meaning that any prediction error relative to these hypotheses will be accorded high expected precision. But this can’t be the proposal, because then the agent would pay attention to all items that aren’t her keys – such items would generate the largest (and hence most precise) prediction errors. It looks like the inverse is needed – the agent must pay attention to the object that generates the least prediction error with respect to the hypothesis ‘this item is my keyset’.

However, this causes the following problem for the PC theory of attention. Recall that precision expectations are expectations of reliable signals. Reliable signals are those that

have a high signal to noise ratio. On the PC theory, signal is just prediction error. So reliable signals are those that generate large prediction errors. For attention to be driven by expectations of precision, then it has to be driven to an area where a large prediction error is generated. However, this is the inverse of what is needed on the presupposition that the relevant hypothesis is ‘this item is my keyset’. Instead, in this case attention is driven to the spatial region that has generated the least prediction error, and so is most accurate – it is driven to the place where one’s keys are located.

Perhaps then the relevant hypothesis ought to be instead ‘it is not the case that this item is my keyset’. This gives us what we need – the largest prediction error will be generated when it does indeed turn out that the item is one’s keyset. One problem with this solution is that it seems ad hoc, vulnerable to the criticism that just about any Bayesian hypothesis can be cooked up to fit the data. Given that negation is a relatively sophisticated concept it is rather implausible that it forms part of the content of our perceptual inference whenever we engage in endogenous cueing tasks. Not only is it a relatively more linguistically sophisticated concept for children to acquire, there is also the issue of how negation is implemented in the predictive coding framework. How might negation be represented in the generative model? What effect does it have on our hypotheses? Is there a vast increase in their complexity, insofar as there are an infinite number of objects that fail to be my keyset and so satisfy my prediction?

(Bowman, Filetti, Wyble, & Olivers, 2013) raise a related worry for the predictive coding account of endogenous attention:

“What makes attention so adaptive is that it can guide towards an object at an unpredictable location – simply on the basis of features. For example, we could ask the reader to find the nearest word printed in bold. Attention will typically shift to one of the headers, and indeed momentarily increase precision there, improving reading. But this makes precision weighting a *consequence* of attending. At least as interesting is the mechanism *enabling* stimulus selection in the first place. The brain has to first deploy attention before a precision advantage can be realized for that deployment” (207, emphasis original).

(Clark, 2013) responds to Bowman et al.’s worry as follows:

“The resolution of this puzzle lies, I suggest, in the potential assignment of precision- weighting at many different levels of the processing hierarchy. Feature-based attention corresponds, intuitively, to increasing the gain on the prediction error units associated with the identity or configuration of a stimulus (e.g. increasing the gain on units responding to the distinctive geometric pattern of a four-leaf clover). Boosting that response (by giving added weight to the relevant kind of sensory prediction error) should enhance detection of that featural cue.

Once the cue is provisionally detected, the subject can fixate the right spatial region, now under conditions of “four-leaf- clover-there” expectation. Residual error is then amplified for that feature at that location, and high confidence in the presence of the four-leaf clover can (if you are lucky!) be obtained” (p.238).

While this answers Bowstein et al.’s worry insofar as one accepts that ‘provisional detection’ can guide spatial attention, it fails to address the original problem raised above because it fails to provide a satisfactory account of how provisional detection is accomplished in the predictive coding framework. To see this, it’s instructive to run through the example of the four-leaf clover more thoroughly. First, the system generates a hypothesis such as: ‘that’s a four-leafed clover’ or ‘that object has four heart-shaped green shapes arranged in a circle.’ The gain on the prediction error units for this hypothesis will be increased. This means that any sensory input that isn’t predicted at any level of the hypothesis generates a large prediction error that will be deemed reliable, and so will be able to preferentially revise the hypothesis. According to Clark, this upping of the gain is enough to enable ‘provisional detection’. But this just leads back to the problem that things that aren’t clovers are going to generate larger prediction errors than things that are, and since the gain has been turned up on any prediction error associated with the hypothesis, such objects will capture our attention preferentially insofar as the system is searching for a clover.

The gist perception solution

At the root of the problem is that the clover hypothesis needs to be selectively applied to the scene (to the space where the clover is located!), but this is exactly what is unknown prior to searching. Lack of spatial expectations for the clover makes it the case that the hypothesis cannot be applied selectively to the scene. However, perhaps provisional detection can be understood along the lines of ‘gist perception’. (Bar, 2003) holds that perception occurs given the system’s ability to first generate a prediction of the ‘gist’ of the scene or object using low spatial frequency visual information that results in a basic- level categorization of the object’s identity (See also Bar et al., 2001; Barrett & Bar, 2009; Oliva & Torralba, 2001; Schyns & Oliva, 1994; Torralba & Oliva, 2003). This then allows for the more fine-grained details to be filled in using the basic-level categorization as a guide. The idea here is that such basic-level categorization could guide selective application of the clover hypothesis, ensuring that it be applied only to objects that have the coarse-grained features of four-leaf clovers. This would then guide attention to the relevant spatial locations, privileging perceptual processing of these areas.

Of course, such a proposal is only a solution if the basic level categorization itself is the result of predictive coding, and here it is unclear as to whether the ‘gist’ is constructed using the hierarchical framework in a purely top-down manner. It certainly does not rely

on high-level hypotheses such as ‘clover’. Constructing the gist of a scene or object would rather be reliant on lower level properties such as shape and color. It is then a further question whether such properties are detected in a feedforward model inconsistent with predictive coding, or predicted in a feedback model consistent with predictive coding (but with lower ‘high level’ hypotheses generating the content of the gist perception). Finally, even if gist perception is amenable to a predictive coding interpretation, there is the further question of exactly how attentional gain fits into the picture here. How, for example, does it prioritize the low spatial frequency information consistent with clover configurations? If the above problem reoccurs for the predictive coding account of gist perception, then it fails as a comprehensive theory of attention.

Moreover, even supposing that a convincing PC model of feature-based attention can be crafted, the account still fails to explain other key aspects of attention, such as emotional salience and non-perceptual attention.

Non-perceptual attention

A complete theory of attention must be able to account for what at least appear to be non-perceptual elements. We can pay attention to our thoughts (sometimes called intellectual attention), ruminations, mind-wanderings, memories and imaginings, and this can occur together with perception – one needn’t close one’s eyes in order to pay attention to one’s thoughts. How can these attentional shifts be accommodated under the present proposal? The picture is rendered even more complicated insofar as sometimes these shifts are exogenous and sometimes they are endogenous – sometimes we decide to pay attention and sometimes our attention is drawn involuntarily inwards. The PC theory of attention is committed to the claim that attention just is the amplification of gain on prediction error, and that this is driven by expected precision. So the proposal would be that we pay attention to our thoughts, imaginings and memories when we expect them to be precise. That is, when they generate a strong signal (large prediction error). What does this amount to? Do we expect our thoughts to be more precise on some occasions rather than others? If so, what learned causal regularity underlies this expectation?

One preliminary suggestion for regulating attention towards our own thoughts might be that it is accomplished via what we term here ‘global gain’. With global gain, learned expected precisions work to either dampen or heighten the bottom up prediction error signal tout court, and as a consequence the gain on the top down hypothesis is heightened or dampened as well in a uniform manner. For example, in extremely poor lighting contexts there is likely to be a significant amount of noise picked up by the visual modality, and so any higher-level hypothesis should give little weight to the prediction errors generated in virtue of the visual modality in this context. Expected precision is the means by which the prediction error generated by the visual modality is dampened down

– in poor lighting contexts bottom-up prediction error is less heavily weighted in virtue of low expected precision from this signal. In such cases, top-down hypotheses are given a higher weight in driving perception, and are unlikely to be significantly revised in light of extremely noisy prediction error. In the extreme version of this case, top down hypotheses aren't modified at all by prediction error. Hohwy (2013) takes this to be an explanation of visual and auditory hallucinations in the face of extremely impoverished sensory input.

The suggestion would then be that we turn our attention inwards when bottom up prediction error is dampened down. Unfortunately, this suggestion is problematic because attention is only identified with the postsynaptic gain on bottom up prediction error units on the PC account, not top down hypotheses. Though such gain may occur, it has no attentional effects. So it cannot be the explanation for internal attention to thought.

A more promising potential avenue for exploration here might be in terms of epistemic emotions (or the emotional content of thoughts more generally, though we leave this out of the discussion for brevity's sake). Epistemic emotions are emotions about the epistemic state of the agent. If one is possession of conflicting evidence, for example p and $\sim p$, then one may feel conflicted or confused. Such conflict may generate a large prediction error given an expectation that one generally does not hold conflicting beliefs. This large prediction error (perhaps felt at the agent level as the epistemic emotion of confusion) may be expected to be precise, given its size. This in turn may draw the agent's attention inwards, towards her own thoughts. That is, such feelings of confusion may be felt before one is aware that one is in possession of such conflicting evidence. So epistemic emotions then serve to guide intellectual attention first by focusing attention inwards and second by sustaining and directing the subsequent searches – one searches for the source of the conflict when one feels confused. While such a suggestion holds promise, one might wonder whether epistemic emotions are really reducible to certain kinds of prediction errors. Moreover, there are substantive issues surrounding the integration of affective salience into the predictive coding model.

Emotional salience

There are many things that are important to our survival and wellbeing that are statistically not very likely to occur in given context. Yet, they can (and ought to) capture our attention in these contexts. This represents a problem for the PC theory of attention, because it is committed to the view that expected precisions are learned statistical regularities, and so one should only pay attention to a given spatial region or object in a context where the signal is expected to be precise. Such cases of emotionally salient objects are counterexamples to the view because they drive attention while nevertheless being unlikely to occur.

For example, suppose you walk your dog every day past a house with an unfriendly Doberman. Though the Doberman is outside in the yard only one out of twenty times, when it is it rushes the fence and startles you. As a consequence, you always attend to the yard when you walk by in order not to be startled. But notice that you don't expect a precise signal – you don't expect the Doberman to be in the yard, because it seldom actually is there. It is rather the extreme unpleasantness of being startled that causes you to attend.

This raises two further potential problems for the PC model of attention. First, it points a more general need to integrate agent-level preferences into the model – not only do we prefer to avoid being startled, but we also have many other preferences that guide attention. Such preferences direct attention in both a top down and bottom up manner – we may notice preferred or highly aversive objects without it being the case that such preferences are relevant to tasks that we are currently engaging in (Niu, Todd, & Anderson, 2012; Todd, Cunningham, Anderson, & Thompson, 2012). The challenge for the PC theory of attention is then to accommodate these additional influences on attention in terms of expected precisions.

Second, the initial example points to the need to factor in the cost of false negatives and false positives, such that standards for expected precision can be adjusted. In an evolutionary context, there is often a significantly higher cost to false negatives over false positives – when an animal's survival is on the line, falsely interpreting noise as signal is the prudentially rational move (within certain boundaries, cf. Stephens, 2001). More colloquially, it's better to be safe than sorry (or dead). On the PC model, attention is driven by signals that are expected to be precise (either because of a bottom up strong signal, or because of a top down expected precision). But attention can also be driven by the cost of getting it wrong – a noisy signal with potentially important information ought to be attended, even when it is not expected to be precise on the PC theory.

Conclusion

In conclusion, while the PC account of endogenous attention works well as an account of endogenous spatial attention, we have argued here that it fails to account for three central features of attention. First, it fails to model endogenous feature-based attention. Second, it does not accommodate non-perceptual forms of attention. Third, it fails to accommodate the influence of affectively salient objects or high cost situations in guiding and capturing attention. While predictive coding provides an attractive account of perception, the account may fail to yield a theory of attention without requiring supplementation that goes beyond the Bayesian framework.

References

- Adams, R. A., Perrinet, L. U., & Friston, K. (2012). Smooth pursuit and visual occlusion: active inference and oculomotor control in schizophrenia. *PLoS One*, 7(10), e47502.
- Bar, M. (2003). A cortical mechanism for triggering top-down facilitation in visual object recognition. *Journal of Cognitive Neuroscience*.
- Bar, M., Tootell, R., Schacter, D. L., Greve, D. N., Fischl, B., Mendola, J. D., ... Dale, A. M. (2001). Cortical Mechanisms Specific to Explicit Visual Object Recognition. *Neuron*, 29(2), 529–535.
- Barrett, L., & Bar, M. (2009). See it with feeling: affective predictions during object perception. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1521), 1325 – 1334.
- Berniker, M., & Kording, K. (2011). Bayesian approaches to sensory integration for motor control. *Wiley Interdisciplinary Reviews: Cognitive Science*, 2(4), 419–428.
- Bishop, C. M. (2006). *Pattern recognition and machine learning*. Springer.
- Bowman, H., Filetti, M., Wyble, B., & Olivers, C. (2013). Attention is more than prediction precision. *The Behavioral and Brain Sciences*, 36(3), 206–8.
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *The Behavioral and Brain Sciences*, 36(3), 181–204.
- Friston, K. J., Daunizeau, J., Kilner, J., & Kiebel, S. J. (2010). Action and behavior: a free-energy formulation. *Biological Cybernetics*, 102(3), 227–60.
- Helmholtz, H. von. (2005). *Treatise on physiological optics*. Mineola: Dover.
- Hobson, J. A., & Friston, K. J. (2012). Waking and dreaming consciousness: neurobiological and functional considerations. *Progress in Neurobiology*, 98(1), 82–98.
- Hohwy, J., Roepstorff, A., & Friston, K. (2008). Predictive coding explains binocular rivalry: an epistemological review. *Cognition*, 108(3), 687 - 701.
- Horga, G., Schatz, K. C., Abi-Dargham, A., & Peterson, B. S. (2014). Deficits in predictive coding underlie hallucinations in schizophrenia. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, 34(24), 8072– 82.
- Huang, Y., & Rao, R. P. N. (2011). Predictive coding. *Wiley Interdisciplinary Reviews: Cognitive Science*, 2(5), 580–593.
- Körding, K. P., & Wolpert, D. M. (2006). Bayesian decision theory in sensorimotor control. *Trends in Cognitive Sciences*, 10(7), 319–26.
- Niu, Y., Todd, R. M., & Anderson, A. K. (2012). Affective salience can reverse the effects of stimulus-driven salience on eye movements in complex scenes. *Frontiers in Psychology*, 3, 336.
- Oliva, A., & Torralba, A. (2001). Modeling the Shape of the Scene: A Holistic Representation of the Spatial Envelope. *International Journal of Computer Vision*, 42(3).
- Rao, R., & Ballard, D. (2004). Probabilistic models of attention based on iconic representations and predictive coding. *Neurobiology of Attention*.

- Schyns, P. G., & Oliva, A. (1994). From blobs to boundary edges: Evidence for time- and spatial-scale-dependent scene recognition. *Psychological Science*, 5(4).
- Seth, A. K., & Critchley, H. D. (2013). Extending predictive processing to the body: emotion as interoceptive inference. *The Behavioral and Brain Sciences*, 36(3), 227–8.
- Seth, A. K., Suzuki, K., & Critchley, H. D. (2011). An interoceptive predictive coding model of conscious presence. *Frontiers in Psychology*, 2, 395.
- Shi, Y. Q., & Sun, H. (1999). *Image and video compression for multimedia engineering*. CRC Press.
- Spratling, M. W. (2008). Predictive coding as a model of biased competition in visual attention. *Vision Research*, 48(12), 1391–408.
- Stefanics, G., Kremlacek, J., & Czigler, I. (2014). Visual mismatch negativity: a predictive coding view. *Frontiers in Human Neuroscience*, 8, 666.
- Stephens, C. (2001). When Is It Selectively Advantageous to Have True Beliefs? Sandwiching the Better Safe than Sorry Argument. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 105(2), 161 – 189.
- Summerfield, C., & Egner, T. (2009). Expectation (and attention) in visual cognition. *Trends in Cognitive Sciences*.
- Todd, R. M., Cunningham, W. A., Anderson, A. K., & Thompson, E. (2012). Affect- biased attention as emotion regulation. *Trends in Cognitive Sciences*, 16(7).
- Torralba, A., & Oliva, A. (2003). Statistics of natural image categories. *Network: Computation in Neural Systems*, 14(3), 391–412.
- Wilkinson, S. (2014). Accounting for the phenomenology and varieties of auditory verbal hallucination within a predictive processing framework. *Consciousness and Cognition*, 30C, 142–155.