Sixty years of predictive perception

Emma K. Ward[1] & Clare Press[1,2]

1. School of Psychological Sciences, Birkbeck, University of London.
2. Wellcome Centre for Human Neuroimaging, UCL

Author note:

This work was funded by a European Research Council (ERC) consolidator grant (101001592) under the European Union's Horizon 2020 research and innovation programme, awarded to CP.

Emma K. Ward ORCID ID: 0000-0003-4295-9167

Conceptualisation: EW & CP
Funding acquisition: CP
Writing, original draft: EW
Writing, review and editing: EW & CP

Sixty years of predictive perception

It is often lamented that scientific progress is slow, perhaps driven by scientists stubbornly persisting with underperforming ideas, testing accounts that are not truly open to the scientific mode of interrogation, and the fundamentally noisy relationship between data and theory (Press, Yon & Heyes, 2022; Oude Maatman, 2021). However, when considering a longer timescale, we should be heartened by the clear evidence that understanding has progressed dramatically. The scientific community understands considerably more than we did several decades ago. Here, we consider the advances in our understanding of the role of prediction in the mind and brain, focusing on its influences on perception. We outline the shifts in explanation, along with methodological developments that have rendered these possible.

Predictive theories of the mind and brain (Clark, 2013; Friston, 2009; Friston, 2010) propose that we integrate incoming sensory evidence with expectations based on our prior experience to form representations of the world. Because sensory input is noisy and the world is generally stable, we can leverage these expectations to bias us toward accurate representations. For example, our vision is less reliable during darkness or fog, so in those conditions we take advantage of what we know of our surroundings. This expectation helps us to conclude that the large indistinct object is a parked car when we're walking down a dark foggy street (see Figure 1, Torralba, 2003). Not only can we *cognitively understand* that the object in this scenario is a car, but we will also *perceive* it as more car-like than the large indistinct blob would allow without these expectations (Clark, 2013)[1].

Some of the knowledge that led to these accounts has been around for decades. For instance, patients with specific disruptions provide a great demonstration of expectations overcoming uncertainty about sensory inputs. Take the description of feature-by-feature object identification, in the case of visual agnosia. Luria (1966, as cited in Thaiss and de Bleser, 1992) described a patient who cannot fluently identify a pair of glasses. However, they could talk through a process in which they recognise a circle, and another circle, joined by a bar, and conclude that this object must be a bicycle. Such disruption in object recognition not only shows potential pathways to resilience in a complex system, but demonstrates clearly that prior experience is essential for perceptual tasks. This explicit, conscious example is of course not the only level at which predictive

[1] Note that this has interesting implications for perception during early development (Ward et al., in prep).

mechanisms operate, but it is a rare opportunity to peer inside the mechanics of the process before the dawn of the technological advances necessary to unpack the underlying mechanisms more precisely.

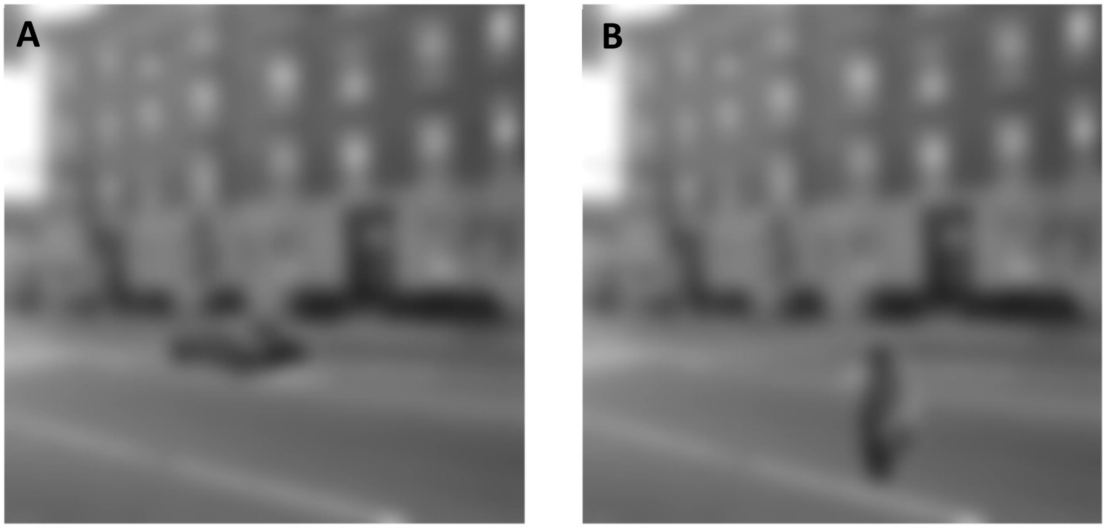

Figure 1. Adapted from Torralba, 2003. Two blurry scenes emulating a dark foggy street, with a prominent object in the foreground of each. In panel A, the object appears to be a car, while in panel B the object appears to be a person.
These objects are in fact the same image superimposed on the background, and only differ due to being rotated by 90 degrees. Our prior experience of the structure of the world scaffolds perception of a car in panel A since cars are the most common horizontal objects of this size in a street scene. Conversely, it generates a percept of a person in panel B, as pedestrians are the most common vertical objects of this size in such scenes.

Although predictive theories of brain and cognition are hundreds of years old (Bubic, von Cramon & Schubotz, 2010), experimental work on predictive motor, sensorimotor, perception and attention processes has rapidly escalated in recent years. Considerable advances in the technology at our disposal, and the ways we conceive of and investigate mental processes, have allowed the field to both collect the data to answer old questions and to generate new questions about predictive theories and their role in perception. The advent of neuroimaging has allowed mechanistic description not possible with patient work, and with it, a cycle of iterative implementation, specification and theorising which allow for the development of new explanatory frameworks (Guest & Martin, 2021). Due to this continuous iterative cycle, theoretical and technological advances are difficult to disentangle, but we have attempted here to outline some of the major advances in our understanding of predictive theories of perception through: 1) the technology used for behavioural experiments, (2) neuroimaging techniques, and (3) shifts in conceptual approaches.

1. Advances via behavioural methods

A host of experimental methods, from adaptation after-effect (Gibson, 1933; Stocker & Simoncelli, 2005) and category learning (Attneave, 1957; Posner & Keele, 1968) to priming (for a review, see Bar, 2004) studies, had already rendered it clear that previous experience directly influences perception not only in the long-term case as demonstrated in visual agnosia patients, but also at timescales in the order of seconds and minutes.

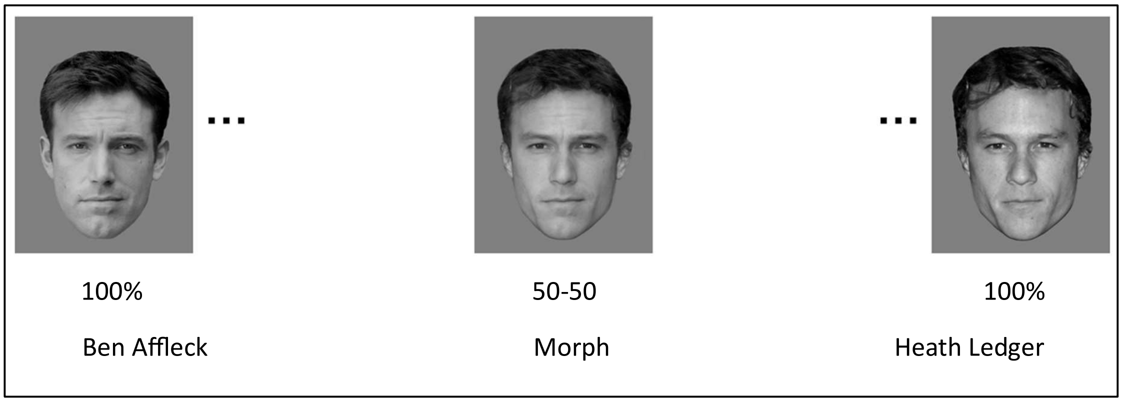

Figure 2. Adapted from Walther, Schweinberger & Kovács, 2013. Adaptation after-effects on facial identity perception.
Familiarise yourself with the three faces, and then look at the face on the left for 5 seconds, followed by looking at the central face. Which person does the central face resemble more? Now look at the face on the right for 5 seconds, followed by the central face. Which person does the central face resemble more now?

Early behavioural paradigms typically relied on participants verbally reporting an embodied experience such as wearing prism glasses (Gibson, 1933) or responding with button presses to static stimuli presented on slides (Posner & Keele, 1968), and while they established a strong evidence base for effects of prior experience on perception, the underlying mechanisms remained somewhat elusive. With the development of stimulus presentation technology, the mystery reduces. For example, with the use of continuous flash suppression (Wolfe, 1984), we have learnt that expected stimuli enter conscious awareness faster (Pinto et al., 2015). With the possibility of fine-tuning of animated avatars, we have learnt that observed expected actions are perceived as higher contrast than unexpected actions (Yon & Press, 2017), while manipulated images have been used to show that expected images are perceived as less blurry (Rossel, Peyrin & Kauffmann, 2023) and higher contrast (Han & van Rullen, 2016) than unexpected images, and virtual reality paradigms have demonstrated that highly expected objects are remembered better than objects which are moderately expected in a given context (Quent, Greve & Henson, 2022). This combination of findings has allowed us to understand that expectations may shape perception by increasing the gain of expected sensory channels relative to unexpected channels (Feuerriegel et al., 2021; Thomas, Rittershofer & Press, 2023), generating increased perception of, sensitivity to, and intensity of percepts that we expect.

2. Advances via neuroimaging techniques

The explosion of fMRI studies in the 1990s (Coltheart, 2006; Price, 2012) enabled spatially-distinct representations in the brain to be distinguished in ways that previously available methods could not, and the development of multivariate approaches (Haynes & Rees, 2006) permitted characterisation of representational fidelity alongside more simple activation profiles. More recently, the application of higher-field fMRI to cognitive neuroscience (Lawrence, Formisano, Muckli & de Lange, 2019) has allowed us to distinguish functional roles of cortical layers. These developments have proven key in our understanding of how the brain makes and transmits predictions, how and when expected and unexpected sensory inputs are represented in the brain, and how these signals are differentiated.

For example, univariate approaches have demonstrated that the brain takes advantage of temporal structure to predict upcoming stimuli (Fiebach & Schubotz, 2006; Kotz, Schwartze & Schmidt-Kassow, 2009), and that when stimuli are expected, mismatches can be processed faster than when they are unexpected (Johnston et al., 2016), perhaps due to sensory cortex pre-activating a template of expected inputs (Kok, Mostert & de Lange, 2017). Furthermore, despite enhanced perception of these expected inputs, they are associated with a reduced signal in early sensory processing regions such as primary visual cortex (Alink et al., 2010; Richter, Ekman & de Lange, 2018; Shergill et al., 2013). Multivariate techniques in the form of linear support vector machines demonstrated that the signal in response to expected stimuli may be weaker, but that pattern classifiers identify them with greater accuracy (Kok, Jehee & de Lange, 2012). It was thus proposed that superior behavioural performance may be due to mechanisms that suppress the representation of unlikely events to generate sharper neural signals. In the last few months, laminar fMRI has demonstrated that these sharper expected signals may exist only in deep cortical structures, and that unexpected events are represented with greater fidelity in superficial layers (Thomas et al., 2023). Thus, mechanisms supporting the role of prediction in perception appear not to consist simply of processes to suppress the gain of unexpected representations. In contrast, predictions may be represented with high fidelity in deep cortical layers, sensory input received in middle layers and superficial layers represent the discrepancy, or error (note that single cell work has also contributed to this mechanistic picture; e.g., Bastos et al., 2020). The error can in turn be passed up the hierarchy for prediction-updating (see Friston, 2005). This mechanistic picture of how exactly visual events are processed required recent developments in neuroimaging and would not have been possible solely with neuropsychological methods and patient studies. In contrast, patient work can show that particular representations may be necessary behaviourally, and in this way, across methods, we progress our understanding of predictive mechanisms.

Integrating these neuroimaging findings with those from behavioural psychophysics, we may thus deduce that perceptual enhancements of expected events are observable over a variety of behavioural measures, and may be due to increased gain of expected representations in deep cortical layers (Aitken et al., 2020; Thomas et al., 2023) that are sometimes activated in advance of the presentation of sensory events (Kok, Mostert & de Lange, 2017). Sensory input signals arrive in middle cortical layers, and an error signal encoding the difference between the expectation and the input is represented in superficial layers. This profile confirms the novel hypotheses generated by predictive theories of perception (Yu et al., 2019) that would have been unthinkable to test 60 years ago.

3. Advances in conceptual approaches

Alongside neuroimaging and modelling now accompanying patient studies as methods for understanding the brain, the last 60 years have seen shifts in broad conceptualisations of mind that have increased the explanatory power of our models. Most notably perhaps, there is a decreased emphasis on modularity and increase in distributed and domain-general theorisation, which has been accompanied by the application of dynamical systems approaches to the mind and brain (Horgan & Tienson, 1992). Such a shift is logical when considering the rapid methods development that was ongoing at the time, and is in line with the postulation that the complexity of a system is not inherent to the system itself but a result of how we are able to observe it (Rosen, 1977). To illustrate the radical change in psychological research following from this shift, take for example the idea of general intelligence. It was assumed up until 20 years ago that perhaps neuropsychology would not concern itself with general intelligence, because they "study bits of the mind and a "bits" (modular) approach is antithetical to the notion" (Anderson, 2005). There was also a community-led discussion on double dissociations, asking whether they were a valid form of enquiry if the brain was not modular (see for example, Dunn & Kirsner, 2003; Gurd & Marshall, 2003; Bullinaria, 2003). The new manners of understanding the brain allow us to understand more precisely the nature of representation emerging through distributed networks between nodes, as well as the content of the representation in nodes themselves.

Perhaps one result of surrendering the view that encapsulated modules are performing distinct operations (Fodor, 1980) is that it is more frequent for discoveries to be taken from one classic domain of cognitive neuroscience to potentially yield insights in another. This has led to both more efficient progress, as fields benefit from knowledge elsewhere, and the emergence of new questions when proposed mechanisms appear to conflict. One clear example comes from the fact that, in contrast with the ideas discussed above that expected information is perceptually enhanced, prevailing in the vision and language sciences, action researchers have claimed that expected input is instead perceptually attenuated (Blakemore, Wolpert & Frith, 1998). For example, these accounts claimed that we cannot tickle ourselves, because we have precise expectations of the incoming sensory input, which therefore generates less intense sensation. This proposed mechanism would allocate fewer resources to processing the expected input to maximise resources for processing surprising information – information that is relevant for model-updating. While it is possible that sensory expectations from acting on the world attenuate perception while other types of expectation enhance it, it is unclear how the brain would implement these opposing influences and why it would be adaptive for the organism to do so. Interestingly, a more domain-general approach has suggested that there is no special influence of expectations on perception in the action domain, as comparable empirical approaches across domains yield comparable, rather than opposite, answers (e.g. Yon, Gilbert, de Lange & Press, 2018; Yon et al., 2021).

The idea that similar underlying mechanisms may support perception across domains, however, does not resolve this conflict, as it does not answer why perception is biased towards expectations in some instances and away in others. These opposing directions of influence have also been shown in other sensory domains. For example, in adaptation after-effects, perception can be strongly and rapidly repulsed from recent experience. In the visual demonstration shown in Figure 2 (from Walther, Schweinberger & Kovacs, 2013), an image of a face representing a 50-50 morph between two identities can be perceived as unambiguously depicting either one of the identities after a short exposure to the opposite one (see Figure 2 caption for demonstration instructions).

It may therefore be the case that the same mechanisms are at work in all sensory modalities, but that there are multiple mechanisms underlying the influence of previous experience on perception (Press et al., 2020a; Feuerriegel, 2023; Teufel & Fletcher, 2020). If the field pursues such explanations, this represents an interesting trajectory, from modular accounts proposing multiple domain-specific mechanisms, to domain-general unified theories, and then subsequent further mechanistic parcellation to accommodate effects that resist a unifying explanation. In this way we hope that science generates progressively better explanations through large-scale shifts and smaller-scale refinement. These advancements necessarily, and importantly, show that the cognitive and neuroscience community is often working collectively to address similar challenges, and that some of these challenges only become apparent when comparing across domains.

This shift from modularity towards emphasis on distributed neural representations has also seen our models largely replace discrete with probabilistic representation. Sixty years ago, cognitive scientists were beginning to understand the importance of variability of prior experience on, for example, learning and categorisation (see for example, Posner & Keele, 1968), but the limitations of the time led to theories based on a one-dimensional statistic. For example, theories of prototypes (Attneave, 1957) and ensemble coding (Ariely, 2001) both posit the average of previously-seen stimuli as embodying the mental representation of a category. The computational limits at the time limited the degrees of experimental freedom, and with these constraints came limits on our ability as scientists to conceive of probabilistic underlying processes. The reduction of these practical constraints has led to considerable theoretical advances, presenting dynamically-generated stimuli, and computing more complex statistics, which further scaffold new empirical discoveries. Describing mental processes as operating over probability distributions rather than discrete point estimates is essential for many proposals within predictive theories (Clark, 2013; Friston, 2009; Friston, 2010; Friston & Stephan, 2007; see Figure 3C). This allows not only for efficient coding of prior expectations as a two-dimensional distribution with a mean and a variance, but also allows for conceptualising more complex cognitive processes than could have been achieved with one-dimensional summaries. Computational proposals have suggested that the expectedness of a stimulus is derived from the extent of overlap between the probability distributions representing the prior expectation and the incoming sensory input, and this overlap has been shown to predict behaviour – such as saccades to informative areas of an image (Itti & Baldi, 2009). Stronger still, recent model-based analyses of fMRI data suggest that not only are these probabilistic formulations useful abstractions, but that the brain may indeed be encoding experiences as probability distributions (van Bergen, Ma, Pratte & Jehee, 2015).

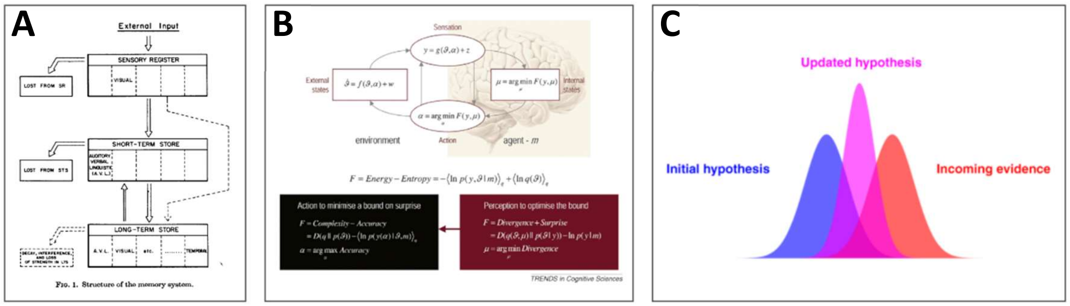

Figure 3. The nuts and bolts of our models have altered considerably across the decades, which have yielded advances in explanatory power. A: Box-and-arrow model of the interactions between sensory input and short- and long-term memory (Atkinson & Shiffrin, 1968). B: Formal model of the interaction of sensory input and prior expectations in the brain (Friston, 2009). C: Schematic of the interaction of sensory input and prior expectations as probability distributions (Yon, Heyes & Press, 2020).

What we do not yet know about expectations and perception

Despite these impressive advances over the past 60 years, there are, of course, still many open questions regarding the influence of expectations on perception. We know that we use our prior experience of the world to predict future observations (Itti & Baldi, 2009; Bubic, von Cramon & Schubotz, 2010; Bar, 2004), but as outlined above, it is still not clear exactly when our perception is biased towards our expectations (Figure 1) and when it is biased away from them (Figure 2). This question is especially important as being biased towards our expectations, on average, increases the chances our perception is veridical in the face of noise, but being biased away from our expectations, on average, increases the chances we perceive an informative new input.

The account that we have been testing states that while perception is initially biased towards expected events, sufficiently unexpected observations trigger reactive processes to generate more precise percepts and allow for accurate model-updating (Press, Kok, Yon 2020a). It was proposed that this could potentially be achieved by phasic noradrenaline release triggered by the unexpected observation, which would increase the gain of sensory input and increase precision of the prediction error (Press, Kok & Yon, 2020a). Under this account, sensory sensitivity to all events may be higher shortly after the unexpected observation. Recent findings, however, have shown that unexpected observations lead to inferior sensory sensitivity across a number of dimensions (Ward & Press, in prep). Another possibility is that unexpected observations lead to poor percepts but trigger either saccades towards the unexpected stimulus (Itti & Baldi, 2005; Press et al., 2020a) or other types of increased epistemic foraging, in which observers seek out more information to reduce uncertainty about the new state of the world (Koenig-Robert et al., preprint; Mirza et al., 2018; Perrykkad, Robinson & Hohwy, 2023; Stahl & Feigenson, 2015). It is therefore possible that a surprising event leads to a weak percept, but also triggers participants to forage for information in the environment to build up a richer percept and learn about the suspected change in the regularities. One yet further possibility is outlined by Feuerriegel (2023), who proposes that instances of perception biased away from previous experience, such as the adaptation after-effect demonstrated in Figure 2, are in fact not governed by expectation mechanisms per se, but purely by neural fatigue. This account still requires bespoke empirical investigation but would likely require surrendering the unifying account of cognition proposed by Friston (2009; 2010) and Clark (2013).

In addition to knowing that expectations do influence our perception, we also know that we need to update those expectations when environmental regularities change, in order to be able to learn (Behrens et al., 2007). For example, when we first see the blurry indistinct object in the dark foggy street, we assume it is a parked car (see Figure 1A). If, however, the object starts to stand up, and we realise that it has legs and appears to be alive, we must quickly change our expectations and our action plan for crossing the street. Nevertheless we do not yet understand how perception is tuned in order to allow efficient learning in these moments of change (Press, Kok & Yon, 2020b), or whether expectations that are learnt as global and context-independent influence perception differently from local, context-dependent expectations (Teufel and Fletcher, 2020). Learning from changes in the environment is a complex task, since we must infer the true nature of the environment from a noisy percept, and in the case of an environmental change we must infer this from a difference in noisy estimates before and after the change. In addition, as outlined above, perception of expected events is often vastly superior to that of surprising events (Bouwer, Honing & Slagter, 2020; Yon & Press, 2017; Rossel, Peyrin & Kauffmann, 2023; Han & van Rullen, 2016; Ward & Press, in prep). It is not yet clear how to reconcile the fact that we form weaker and less precise percepts of surprising events with the fact that we do successfully learn from changes in the environment that these surprising events indicate.

A final empirical datapoint that may help to solve the mystery is that, despite this paucity of the percept when the input is unexpected, unexpected observations are more accurately decoded from neural signals at particular moments in time (EEG; Rittershofer et al., in prep) and from superficial layers of visual cortex (laminar fMRI; Thomas et al., 2023). This combination of perceptual and neural findings raises the possibility that model updating proceeds alongside poor conscious perception. This possibility has been explored in a recently-proposed framework (Soto, Sheikh & Rosenthal, 2019), although predictive theories cannot currently account for this empirical pattern.

The fast pace of methodological and theoretical developments will hopefully allow us to answer these questions within the next few years.

Conclusion

The combined approach of using imaging in combination with patient work, along with a variety of stimulus presentation and computational modelling developments, has allowed a step change across 60 years in understanding how our prior experience of the structure of the world shapes our representations of it. There may be many disputes in the domain of prediction but there are arguably more agreements than disagreements. A majority of scientists now believe that sensory and higher order processing proceeds via integrating what we already knew with the sensorium in the here-and-now. Understanding will continue to advance as we develop further an array of methods to compare conclusions in lab-based and naturalistic settings (Vigliocco et al., 2023) and think long and hard about explanations that fit the complexity of the empirical picture (Press, Yon & Heyes, 2022). We are excited to see what the next 60 years will bring.

References

Aitken, F., Menelaou, G., Warrington, O., Koolschijn, R. S., Corbin, N., Callaghan, M. F., & Kok, P. (2020). Prior expectations evoke stimulus-specific activity in the deep layers of the primary visual cortex. *PLoS Biology, 18*(12), e3001023.

Alink, A., Schwiedrzik, C. M., Kohler, A., Singer, W. & Muckli, L. (2010). Stimulus predictability reduces responses in primary visual cortex. *Journal of Neuroscience,* 30, 2960–2966.

Anderson, M. (2005). Cortex forum on the concept of general intelligence in neuropsychology. *Cortex, 41*(2), 99-100.

Ariely, D. (2001). Seeing sets: Representation by statistical properties. *Psychological Science, 12*(2), 157-162. https://doi.org/10.1111/1467-9280.00327

Atkinson, R. C., & Shiffrin, R. M. (1968). Human memory: A proposed system and its control processes. In *Psychology of Learning and Motivation* (Vol. 2, pp. 89-195). Academic Press.

Attneave, F. (1957). Transfer of experience with a class-schema to identification-learning of patterns and shapes. *Journal of Experimental Psychology, 54*(2), 81.

Bar, M. (2004). Visual objects in context. *Nature Reviews Neuroscience, 5,* 617–629.

Bastos, A.M., Lundqvist, M., Waite, A.S., Kopell, N., & Miller, E.K. (2020). Layer and rhythm specificity for predictive routing. *Proceedings of the National Academy of Sciences USA, 117,* 31459-31469.

Bechtel, W. (2009). Looking down, around, and up: Mechanistic explanation in psychology. *Philosophical Psychology, 22*(5), 543-564.

Behrens, T. E., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience, 10*(9), 1214-1221.

Blakemore, S. J., Wolpert, D. M., & Frith, C. D. (1998). Central cancellation of self-produced tickle sensation. *Nature Neuroscience, 1*(7), 635-640.

Bouwer, F. L., Honing, H., & Slagter, H. A. (2020). Beat-based and memory-based temporal expectations in rhythm: similar perceptual effects, different underlying mechanisms. *Journal of Cognitive Neuroscience, 32*(7), 1221-1241.

Bubic, A., Von Cramon, D. Y., & Schubotz, R. I. (2010). Prediction, cognition and the brain. *Frontiers in Human Neuroscience, 4*, 1094.

Bullinaria, J. A. (2003). Dissociation in connectionist systems. *Cortex, 39*(1), 142-144.

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. Behavioral and Brain Sciences, 36(3), 181-204. https://doi.org/10.1017/S0140525X12000477

Coltheart, M. (2006). What has functional neuroimaging told us about the mind (so far)? *Cortex, 42*(3), 323-331. https://doi.org/10.1016/s0010-9452(08)70358-7

Dunn, J. C., & Kirsner, K. (2003). What can we infer from double dissociations? *Cortex, 39*(1), 1-8.

Feuerriegel, D. (2023). Adaptation in the Visual System: Networked Fatigue or Suppressed Prediction Error Signalling? *PsyArXiv,* https://osf.io/preprints/psyarxiv/73hb4/

Feuerriegel, D., Yook, J., Quek, G. L., Hogendoorn, H., & Bode, S. (2021). Visual mismatch responses index surprise signalling but not expectation suppression. *Cortex, 134*, 16-29. https://doi.org/10.1016/j.cortex.2020.10.006

Fiebach, C. J., and Schubotz, R. I. (2006). Dynamic anticipatory processing of hierarchical sequential events: a common role for Broca's area and ventral premotor cortex across domains? *Cortex 42*, 499–502.

Fodor, J. A. (1980). Methodological solipsism considered as a research strategy in cognitive psychology. *Behavioral and Brain Sciences, 3*, 63-109.

Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society B: Biological sciences*, 360(1456), 815-836.

Friston, K. (2009). The free-energy principle: a rough guide to the brain? *Trends in Cognitive Sciences, 13*(7), 293-301.

Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience, 11*(2), 127-138.

Friston, K. (2018). Does predictive coding have a future? *Nature Neuroscience, 21*, 1019–1021. https://doi.org/10.1038/s41593-018-0200-7

Friston, K. J., & Stephan, K. E. (2007). Free-energy and the brain. *Synthese, 159*, 417-458.

Gibson, J. J. (1933). Adaptation, after-effect and contrast in the perception of curved lines. *Journal of Experimental Psychology, 16*(1), 1.

Greve, A., Cooper, E., Tibon, R., & Henson, R. N. (2019). Knowledge is power: Prior knowledge aids memory for both congruent and incongruent events, but in different ways. *Journal of Experimental Psychology: General, 148*(2), 325.

Guest, O., & Martin, A. E. (2021). How Computational Modeling Can Force Theory Building in Psychological Science. *Perspectives on Psychological Science, 16*(4), 789–802. https://doi.org/10.1177/1745691620970585

Gurd, J. M., & Marshall, J. C. (2003). Dissociations: double or quits? *Cortex, 39*(1), 192-195.

Han, B. & VanRullen, R. Shape perception enhances perceived contrast: evidence for excitatory predictive feedback? *Scientific Reports, 6,* 22944 (2016).

Haynes, J.D. & Rees, G. (2006). Decoding mental states from brain activity in humans. *Nature Reviews Neuroscience, 7*, 523–534. https://doi.org/10.1038/nrn1931

Horgan, T., & Tienson, J. (1992). Cognitive systems as dynamical systems. *Topoi, 11*, 27–43.

Itti, L., & Baldi, P. (2009). Bayesian surprise attracts human attention. *Vision Research, 49*(10), 1295-1306.

Johnston, P., Overell, A., Kaufman, J., Robinson, J., & Young, A. W. (2016). Expectations about person identity modulate the face-sensitive N170. *Cortex, 85*, 54-64.

Koenig-Robert, R., Pace, T., Pearson, J., & Hohwy, J. (2023). Time-resolved hierarchical frequency-tagging reveals markers of predictive processing in the action-perception loop. *Authorea Preprints.* https://doi.org/10.22541/au.167338513.30310833/v1

Kok, P., Jehee, J. F., & De Lange, F. P. (2012). Less is more: expectation sharpens representations in the primary visual cortex. *Neuron, 75(2),* 265-270.

Kok, P., Mostert, P. & de Lange, F. P. (2017). Prior expectations induce prestimulus sensory templates. *PNAS, 114*, 10473–10478.

Kotz, S. A., Schwartze, M., and Schmidt-Kassow, M. (2009). Non-motor basal ganglia functions: a review and proposal for a model of sensory predictability in auditory language perception. *Cortex, 45*, 982–990.

Kwisthout, J., Bekkering, H., & Van Rooij, I. (2017). To be precise, the details don't matter: on predictive processing, precision, and level of detail of predictions. *Brain and Cognition, 112*, 84-91.

Lawrence, S. J., Formisano, E., Muckli, L., & de Lange, F. P. (2019). Laminar fMRI: Applications for cognitive neuroscience. *Neuroimage, 197*, 785-791.

Leube, D. T. et al. The neural correlates of perceiving one's own movements. *Neuroimage* 20, 2084–2090 (2003).

Lively, S. E., Logan, J. S., & Pisoni, D. B. (1993). Training Japanese listeners to identify English /r/ and /l/.II: The role of phonetic environment and talker variability in learning new perceptual categories. *The Journal of the Acoustical Society of America*, 94 (3 Pt 1), 1242 – 1255. https://doi.org/10.1121/1.408177

Luria, A.R. (1966). Higher Cortical Functions in Man. London: Tavistock Publishing.

Mirza, M. B., Adams, R. A., Mathys, C., & Friston, K. J. (2018). Human visual exploration reduces uncertainty about the sensed world. *PloS one, 13*(1), e0190429.

Oude Maatman, F. (2021). Psychology's theory crisis, and why formal modelling cannot solve it. https://psyarxiv.com/puqvs/

Perrykkad, K., Robinson, J. E., & Hohwy, J. (2023). Foraging for the self: Environment selection for agency inference. *Psychonomic Bulletin & Review, 30*(2), 608-620.

Pinto, Y., van Gaal, S., de Lange, F. P., Lamme, V. A. F. & Seth, A. K. (2015). Expectations accelerate entry of visual stimuli into awareness. *Journal of Vision, 15*, 13.

Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology, 77*(3, Pt.1), 353–363. https://doi.org/10.1037/h0025953

Press, C., Kok, P., & Yon, D. (2020a). The perceptual prediction paradox. *Trends in Cognitive Sciences, 24(1),* 13-24.

Press, C., Kok, P., & Yon, D. (2020b). Learning to perceive and perceiving to learn. *Trends in Cognitive Sciences, 24*(4), 260-261.

Press, C., Yon, D., & Heyes, C. (2022). Building better theories. *Current Biology, 32(1),* R13-R17.

Price, C. J. (2012). A review and synthesis of the first 20 years of PET and fMRI studies of heard speech, spoken language and reading. *Neuroimage, 62*(2), 816-847. https://doi.org/10.1016/j.neuroimage.2012.04.062

Quent, J. A., Greve, A., & Henson, R. N. (2022). Shape of U: The Nonmonotonic Relationship Between Object–Location Memory and Expectedness. *Psychological Science, 33*(12), 2084-2097. https://doi.org/10.1177/09567976221109134

Richter, D., Ekman, M., & de Lange, F. P. (2018). Suppressed sensory response to predictable object stimuli throughout the ventral visual stream. *Journal of Neuroscience, 38*(34), 7452-7461.

Rosen, R. Complexity as a system property. *International Journal of General Systems, 3*(4), 227-232.

Rossel, P., Peyrin, C., & Kauffmann, L. (2023). Subjective perception of objects depends on the interaction between the validity of context-based expectations and signal reliability. *Vision Research, 206*, 108191.

Shergill, S. S., White, T. P., Joyce, D. W., Bays, P. M., Wolpert, D. M., & Frith, C. D. (2013). Modulation of somatosensory processing by action. *Neuroimage, 70*, 356-362.

Soto, D., Sheikh, U. A., & Rosenthal, C. R. (2019). A novel framework for unconscious processing. *Trends in Cognitive Sciences, 23*(5), 372-376. https://doi.org/10.1016/j.tics.2019.03.002

Stahl, A. E., & Feigenson, L. (2015). Observing the unexpected enhances infants' learning and exploration. *Science, 348*(6230), 91-94. https://doi.org/10.1126/science.aaa3799

Stocker, A. A., & Simoncelli, E. (2005). Sensory adaptation within a Bayesian framework for perception. *Advances in Neural Information Processing Systems,* 18.

Teufel, C., & Fletcher, P. C. (2020). Forms of prediction in the nervous system. *Nature Reviews Neuroscience*, 21(4), 231-242.

Thomas, E. R., Haarsma, J., Nicholson, J., Yon, D., Kok, P., & Press, C. (2023). Predictions and errors are distinctly represented across V1 layers. bioRxiv, 2023-07. https://www.biorxiv.org/content/10.1101/2023.07.11.548408v1.abstract

Torralba, A. (2003). Contextual Priming for Object Detection. *International Journal of Computer Vision, 53*, 169–191. https://doi.org/10.1023/A:1023052124951

van Bergen, R. S., Ji Ma, W., Pratte, M. S., & Jehee, J. F. (2015). Sensory uncertainty decoded from visual cortex predicts behavior. *Nature Neuroscience, 18*(12), 1728-1730.

Vigliocco, G., Convertino, L., De Felice, S., Gregorians, L., Kewenig, V., Mueller, M. A., ... & Spiers, H. (2023). Ecological Brain: Reframing the Study of Human Behaviour and Cognition. https://psyarxiv.com/zr4nm/

Walther, C., Schweinberger, S. R., & Kovács, G. (2013). Adaptor identity modulates adaptation effects in familiar face identification and their neural correlates. *PloS one, 8*(8), e70525.

Ward, E. & Press, C. (in prep). Surprise impairs perception of surprising, and incidental, events.

Ward, E., Rutar, D., Zaadnoordijk, L., Poli, F. & Hunnius, S. (in prep). Beyond the Adult Mind: A Developmental Framework for Predictive Processing in Infancy.

Wolfe, J. M. (1984). Reversing ocular dominance and suppression in a single flash. *Vision Research, 24*, 471-478

Yon, D., Gilbert, S.J., de Lange, F.P. & Press, C. (2018). Action sharpens sensory representations of expected outcomes. *Nature Communications* 9, 4288. https://doi.org/10.1038/s41467-018-06752-7

Yon, D., Heyes, C. & Press, C. (2020). Beliefs and desires in the predictive brain. *Nature Communications, 11*, 4404 (2020). https://doi.org/10.1038/s41467-020-18332-9

Yon, D., & Press, C. (2017). Predicted action consequences are perceptually facilitated before cancellation. *Journal of Experimental Psychology: Human Perception and Performance, 43*(6), 1073.

Yon, D., Zainzinger, V., de Lange, F. P., Eimer, M., & Press, C. (2021). Action biases perceptual decisions toward expected outcomes. *Journal of Experimental Psychology: General, 150*(6), 1225. http://dx.doi.org/10.1037/xge0000826

Yu, Y., Huber, L., Yang, J., Jangraw, D. C., Handwerker, D. A., Molfese, P. J., ... & Bandettini, P. A. (2019). Layer-specific activation of sensory input and predictive feedback in the human primary somatosensory cortex. *Science Advances, 5*(5), eaav9053.