Title

Bayesian learning models of pain - a call to action

Authors

*Abby Tabor[1], Christopher Burr[2]

[1]Centre for Pain Research, Department for Health, University of Bath, Claverton Down Rd, Bath, BA2 7AY. a.tabor@bath.ac.uk

[2]Department of Computer Science, University of Bristol, Merchant Venturers Building, Woodland Road, Clifton BS8 1UB. chris.burr@bristol.ac.uk

* Denotes corresponding author

Abstract

Learning is fundamentally about action, enabling the successful navigation of a changing and uncertain environment. The experience of pain is central to this process, indicating the need for a change in action so as to mitigate potential threat to bodily integrity. This review considers the application of Bayesian models of learning in pain, which inherently accommodate uncertainty and action, which, we shall propose are essential in understanding learning in both acute and persistent cases of pain.

Highlights

- The experience of pain sits awkwardly in traditional stimulus-response paradigms
- Accommodating uncertainty and action is imperative to learning models of pain
- Bayesian models provide a normative, probabilistic account of learning in pain
- Learning in pain is conceptualised as an ongoing prediction of the consequences of action

Introduction

The process of learning is fundamentally about action. In order to successfully navigate our environment, we must continually learn about the ever-changing limits of our body and the constraints that it imposes upon our interaction with the world. The experience of pain is central to this process, indicating the point at which our bodily integrity is potentially compromised through action.

The interaction between pain and learning can be better understood from an evolutionary perspective, by adopting the concept of the explore-exploit dilemma [1]. When our bodily integrity is threatened, we typically withdraw or rest (exploit) to allow sufficient recovery to within bodily limits, at which point we decide to interact (explore) within our niche. We learn over time when it is best to exploit and when to explore in order to promote adaptive behaviour [2,3].

Learning in pain, however, is not straightforward, owing to the complexity that comprises bodily integrity and worldly state. As a consequence, we find ourselves confronted with the reality that in some cases pain persists, seemingly decoupled from acute protection and adaptive behaviour. This necessarily goes beyond responding to and learning about a nociceptive signal, extending to an overall appraisal of the bodily and sociocultural environments in which we exist [4,5]. Adequately accounting for such a rich and diverse set of interactions is the challenge faced in establishing a learning model in pain.

Current application of learning models in pain

Over the last 40 years, associative learning models have come to dominate our conception of learning in pain [6]. These accounts are pervasive in different forms across the pain field, from Pavlovian (habitual) to Operant (instrumental) conditioning in behavioural psychology [7,8], extending to reinforcement learning and temporal difference models in computational neuroscience [9–12]. Operationalised through the Rescorla-Wagner model, the heart of associative learning models lies in the concept of an associative weight between stimulus and response, ranging from immediate, reflexive stimulus-response (model-free) to more complicated goal-directed actions, which alter proceeding stimuli (model-based) [13].

Through the application of associative learning theory, it is posited that persistent pain reflects the generalisation of pain-related responses and maintained avoidance behaviour [8,17]. This conceptualisation has shaped our understanding of pain in the behavioural sciences, an influence seen from scientific investigation to clinical management.

Yet, the *experience of pain* sits awkwardly in these traditional stimulus-response models [21,22]. In light of recent advances across neuroscience and behavioural domains, there is a growing consensus that perceptual experience is a predictive process, in which learners actively seek information to update their prediction of their internal and external environment [23,24]. This is problematic for traditional associative learning models in pain for several reasons. Firstly, pain is classically posited as a stimulus and conflated with nociception, which downplays the significance of *pain as an experience* and its explanatory role within theories of learning. Secondly, traditional associative learning models the state of the learner as a series of punctate values at any given time [26], which belies the learner's uncertainty [25–28]. Finally, associative models do not adequately accommodate the active nature of the learner (i.e. being able to actively explore and intervene in their environment) [29]. It is proposed that these challenges for traditional learning models may be overcome by taking a Bayesian approach to learning in pain.

The Bayesian Framework

Bayesian approaches to cognition comprise many distinct models and theories, used in a variety of domains, and spanning distinct levels of explanation. Often, these distinct approaches are grouped under the label 'Bayesian Brain hypothesis' [30,31], despite their many differences. This review will focus on the underlying Bayesian model that informs these approaches, specifying the Bayesian derivative where appropriate.

To date, the application of Bayesian models in pain has been limited to the description of perceptual experience, presenting pain as part of a probabilistic inference process that is shaped through the optimal integration of informative cues [27]. These models propose a mechanism for determining the hidden (latent) causes of encountered sensory information, summarised in a generative model [32]. In Bayesian terms, this is achieved through the weighted integration of prior experience and current (potentially multisensory) information, represented using probability distributions that reflect the agent's subjective uncertainty—the optimal integration of these probability distributions is given by Bayes' rule [33] (Fig. 1).

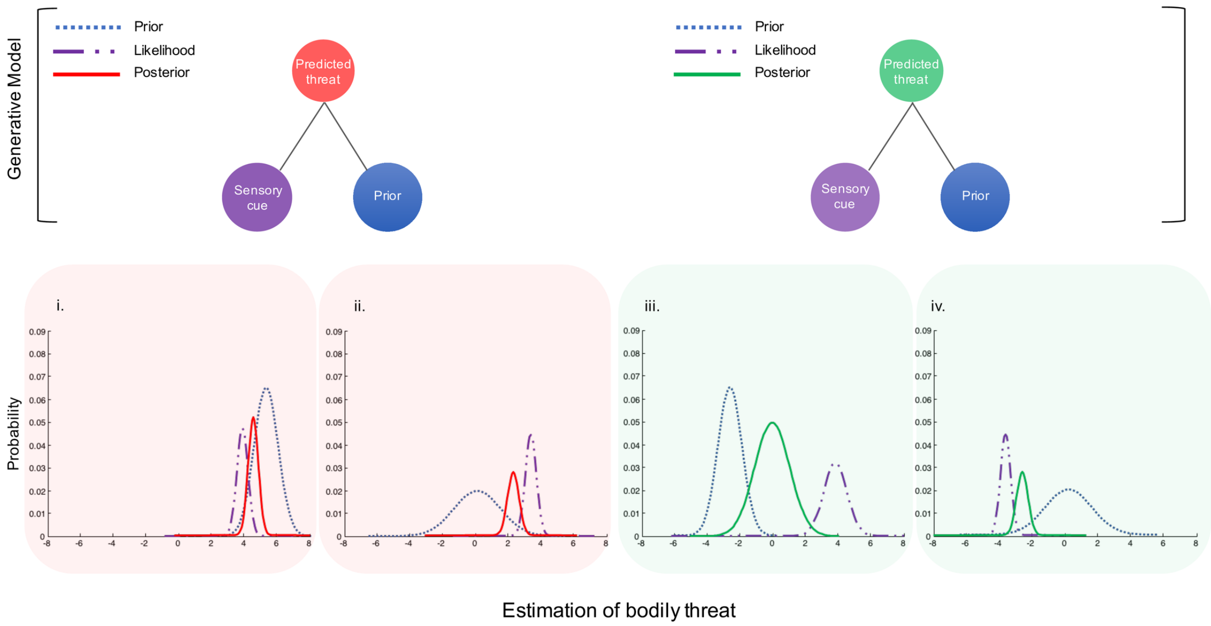

***Fig. 1. Generative models: Prediction of bodily threat (i-iv.)**. A generative model provides the framework from which predictions of the hidden causes of sensory consequences are generated (posterior), these are continually informed by multisensory sensory cues (likelihood) and previous encounters (prior). The relative precision, reflected in the probability density of these elements, influences the prediction. The more precise (narrow probability distribution), the greater the influence on the prediction. Threat panels (Left: i-ii) demonstrate the relative contribution of either a relatively precise prior (i) or precise likelihood (ii), the resultant prediction of threat is drawn toward the more precise source of information. In these cases, the sensory cue (likelihood) is the same in both panels, yet the relative precision of the prior determines the overall prediction of threat. Safety Panels (Right: iii-iv) demonstrate how the same relative precision can influence the prediction of negative threat, or safety. A precise prior, even in the presence of objective threat-based sensory cues, can influence the overall prediction to reflect safety (placebo effect) (iii). In contrast, an imprecise prior has less influence on the posterior (negative threat/safety) (iv). These hypothetical generative models demonstrate the possible decoupling of objective sensory information from experience, by accounting for the precision of the prior, which reflects the ongoing learning of the individual in keeping with previous experiences, homeostatic bounds and sociocultural constraints.*

Although not directly about learning, these accounts expose the fundamental elements of the Bayesian approach: a generative model, subjective uncertainty, and variable precision-weighting. It is through the inherent encoding of the learner's uncertainty that Bayesian models can shift away from specific associative weighting between variables towards a learning account that is both predictive and active. This is a significant theoretical development [26], which will form the basis of the proceeding review.

Learning under uncertainty

In Bayesian approaches, learners are assumed to have only indirect access to the state of their internal and external environment and must, therefore, infer their values on the basis of ambiguous and often incomplete information [34]. In contrast to associative learning models, Bayesian models encode uncertain beliefs about the world as probability distributions [35]. They assume that learners maintain multiple hypotheses (with differing degrees of belief) that reflect a range of candidate predictions about the state of the body and the world. This invokes the notion of a generative model (Fig.1), which can be used to *generate* the expected sensory consequences that may arise from hidden (latent) states of the environment, and in absence of external stimulation [36,37].

According to Bayesian models, learning occurs through the adjustment of the prior distribution (e.g. estimated threat), according to Bayes rule, when new sensory cues are encountered. This asserts that over time a learner attempts to predict, with increasing finesse, the state of the world. Rather than veridical reflections, these predictions are an integration of probability distributions pertaining to the precision of the information.

An emerging framework, derived from a Bayesian approach, known as predictive processing [23,38–40] casts the inferential process in probabilistic modelling as a matter of *prediction-error minimisation*. According to this view, the learner's generative model gives rise to multiple top-down predictions that are met by incoming sensory information (prediction error). This is a competitive process, where the prediction that best captures the incoming sensory information is selected, and perception arises as a result of successful prediction-error minimisation[1].

The concept of prediction error here represents predictions and prediction-errors as probability distributions, thus retaining the inherent encoding of uncertainty of an agent's beliefs that is common to Bayesian approaches. In predictive processing, specifically, this uncertainty is managed by precision-weighting mechanisms, which modulate the variance associated with the respective distributions, in order to contribute to the overall goal of minimising prediction error [41]. From this perspective, the learner's principle motivation is to minimise the discrepancy between their prediction of the world and the sensory consequences of it (prediction error), in order to ensure they maintain an accurate model of their world (Fig.2). At a cortical level, it has been proposed that precision weighting of prediction errors is mediated by dopamine, with the potential to influence both accurate and aberrant learning [42–44].

As we shall explore, the learner can minimise prediction error in two ways: by updating the parameters of their generative model in order to better predict the future sensory consequences of action, or by holding the model fixed and altering their action within the world to sample information that better reflects their predictions. These mechanisms are described under the *Active Inference* framework [45].

[1] For a non-technical, conceptual introduction to the Predictive Processing framework see [23]. For for an overview of how the free-energy principle applies to the brain see [57].

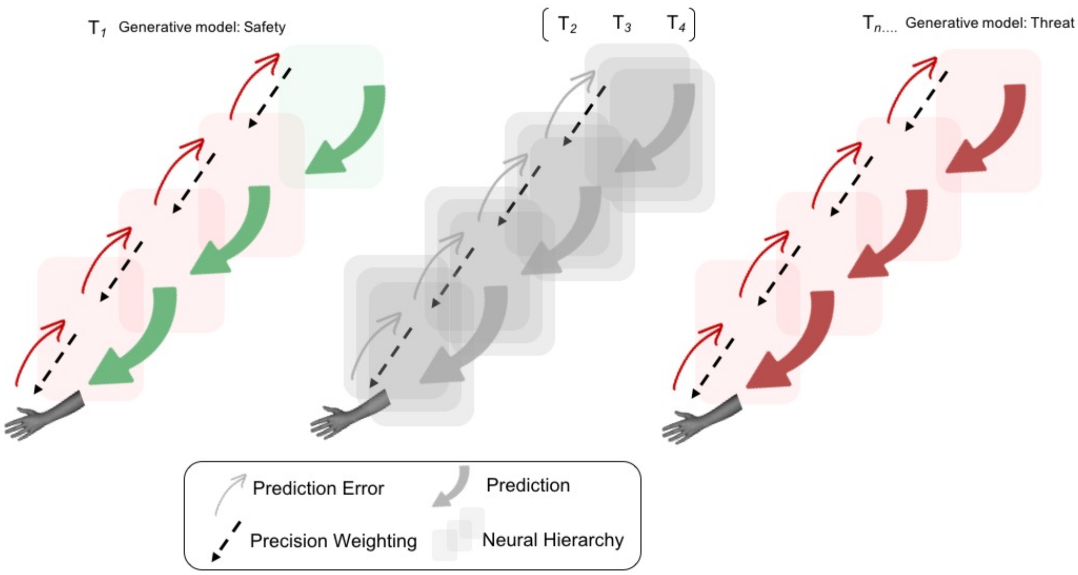

***Fig.2. Hierarchical Predictive Processing: from safety to threat**. Proposed within a neural hierarchy, generative models are shaped over time to reflect the precision weighting of information. There is a continual bidirectional flow of information at each level of the neural hierarchy involving top-down predictions, prediction error, and the precision-weighting of prediction error. Schematically represented over time, an initial generative model encompassing bodily safety entails the prediction of low bodily threat as a consequence of action. Over time, in the presence of prediction error (a deviation from predicted bodily safety or predicted bodily threat), the generative model is updated to reflect an alteration in action consequences, that of threat. It is suggested that the ability to flexibly update this prediction of threat, in the presence of new sensory evidence (e.g. safety cues), is imperative to the resolution of the need to experience pain.*

Active Learning

The inherent uncertainty encoded in the agent's probability distributions not only satisfies learning paradigms that are typically challenging for associative theories (e.g backward blocking; see [26]), it crucially affords the agent an active role in reducing uncertainty. Active learning under this formulation is not simply the provision of an adequate sample space (spatial and temporal), it rests on the crucial ability of the learner to intervene in their world, sculpting the sensory consequences of their actions according to what is deemed most salient. The consequences of the learner's actions can either support or disconfirm the predictions of the consequences of action, offering multiple means by which to reduce uncertainty [46,47]. These considerations of active learning recognise ecological validity from the perspective of being in, and acting upon, the world, and where actions are taken based on the ongoing (motivational) homeostatic drives of the biological agent.

Active Inference[2]—a component of the predictive processing framework—extends these basic commitments and transforms the role of the learner in pain, from a passive processor of information, to a dynamic predictor of the relationship between the external and internal world. A key claim of the active inference model is that embodied action occurs as a result of

[2] For a review of active inference, which casts it as a process of descending projections (predictions) from motor cortex, see [43]. Other accounts have implicated the dopaminergic system as playing a key role in active decision-making, while also casting this within the framework of ecological psychology [44]. And, more recently, the active inference framework has been extended to incorporate homeostatic control [40, 51]. For a less technical overview, including empirical and theoretical support, see chapter 4 of [23].

an agent predicting (inferring) the outcomes of certain policies (e.g. reaching for a cup), along with their associated precision estimations. The process of predicting future consequences of actions (i.e. associated sensory information) leads to overt behaviour through the activation of classical reflex arcs by downwards projections from motor cortex [46]. An illustrative example would be a policy that controls an agent's task of reaching for a cup. Prior to enacting the reaching behaviour, the predictions associated with grasping the cup will be unfulfilled, and therefore result in error signals. However, instead of updating the generative model, the agent can instead take actions that lead to the fulfilment of error signals by actually reaching to grab the cup. In cases where the agent predicts that a certain policy will also likely lead to the experience of pain (e.g. bending down to pick up a heavy box), the agent may be reluctant to enact the respective behaviour, or choose to avoid it altogether.

Seth [48] and others [51] have extended the active inference framework to account for autonomic regulation, arguing that similar predictions generated by the AIC are sent to the autonomic system via smooth muscles to activate autonomic reflexes in a similar manner as earlier described in the case of proprioception. By focussing on the embodied nature of the agent, active inference creates an intuitive segue that unites learning about the state of external world (exteroception) with the state of the internal world (interoception). The same predictive mechanisms that are responsible for predicting sensory states of the external environment are also responsible for regulating the internal environment [48–50] and for providing additional sources of information related to motivational drives [51]. Although often separated in traditional theories, perception and action are entwined in active inference, due to their dual-role in minimising uncertainty [15].

The proposition that a single underlying mechanism (i.e. precision-weighted prediction-error minimisation) underlies learning about the condition of the body, has provided instrumental guidance for describing the generation of aberrant bodily predictions and the development of persistent pathological conditions [42,43,52–54]. It is suggested that persistent pain can be formulated in such a way [55]. To illustrate this, the experience of pain is mapped onto the 'warning light' scenario, proposed by Adams et al, 2013:

> *Consider a circumstance in which you are experiencing knee pain, you predict, with high precision, that the consequences of your action in the world will compromise the integrity of your body. Minor fluctuations in your interoceptive sensory cues (prediction errors) are assigned high precision, which serve to confirm the prediction of potential threat and propagate your experience of pain. You decide to visit your doctor who is unable to determine a specific cause for your pain, they even present you with your x-ray that shows "no structural cause for your pain". Your first thought is that your doctor has missed something, that there must be something else going on, or that the x-ray has been misinterpreted. From your perspective all of these are plausible hypotheses that accommodate the evidence that is available to you. However, from the doctor's perspective, without the knowledge that informs your prediction of bodily threat, your suspicions seem irrational.*

This adapted account highlights the consequences of precision weighting of information in the experience of pain. What is suggested is a decoupling between sensory input and subjective experience, where the latter is dependent on the relative precisions afforded to predictions and prediction error (Fig.2). The learner in pain updates the precision weighting of information that reflects their generative model in a changing world, informing whether to exploit or explore[3]. This places experiences of the body, whether well-defined through

---

[3] Some have proposed that precision-weighting may also be responsible for the transient switching between online and offline control [41]—allowing an agent to deliberate about some future policy, prior to taking action within the world. Although generative models play a central role in guiding online

disease process or medically unexplained, on a continuum [55]. What distinguishes them is the accuracy with which they account for the underlying physiological condition of the body.

Persistent pain, from this view, occurs as aconsequence of precision: either via a precise prediction of bodily threat (top-down) or through aberrant precision weighting of sensory information (bottom up). In both cases, the prediction of bodily threat persists, and so with it the experience of pain, detached from veridical evidence of tissue damage and unchallenged by information assigned less precision. Altering the experience of pain this lies in the ability to promote the flexible reassignment of precision weighting, which in turn alters the individual's prediction about their body and the world.

The description of learning in pain thus becomes one that concerns optimal precision weighting over time. Importantly, under normative models, optimality does not pertain to accuracy. As such, aberrant but precise predictions of bodily threat (e.g high precision-weighting of noisy sensory signals), and an accompanying experience of pain, may persist in the absence of an objective reality of threat. No more or less real, all experiences of the body are a reflection of our evolutionary history, sociocultural present and action-oriented future.

Discussion

One core pursuit of learning models in pain is to adequately accommodate the phenomena of acute and persistent cases. That is, why do the majority of people experience pain as transitory—an experience that efficiently promotes acute withdrawal, mitigating further harm—while a significant minority continue to experience pain in a way that seemingly contravenes optimal behaviour?

We have broadly considered Bayesian models and their relevance to learning in pain. It is proposed that in order to accommodate the ecological validity of the learner in pain, the concepts of uncertainty and active learning must be addressed. As such, derivatives of the Bayesian model have been presented, which attempt to re-conceptualise the learner as an action-oriented predictor of their environment.

An advantage of this approach is that learning in pain is considered under a unifying framework. The experience of pain becomes a problem of precision-weighting, inherently contextualised in relation to previous experience and future endeavour; both the resolution and persistence of pain lies within one's ability to continually update the predictions of bodily state.

The approaches that are described are not wholly opposed to the concepts present in associative learning accounts (e.g. kalman filter and the Rescorla Wagner model) [26]. However, a probabilistic formulation of learning promotes an account that naturally extends to the body and action [56], and is highly relevant to learning in pain, whereby the active sampling of one's environment is fundamentally altered.

---

behaviour (i.e. active inference), by decoupling generative models from the incoming stream of sensory information (prediction errors), through the use of selective modulation of incoming prediction errors (precision-weighting), generative models may also guide deliberative processes such as planning and offline reasoning [41,51]. This flexible switching between offline and online control could be viewed as a type of arbitration mechanism for model-free and model-based forms of behavioural control, albeit one that may be best viewed as more of a continuum of cases, rather than a well-delineated set of options [41].

Bayesian formulations have proffered much, not least a unifying theory of mind [57]. Yet, with such promise comes inevitable pitfalls [58–60], a number of which require consideration here.

This review has focussed predominantly on the implementation of such models at an instrumental level, describing the macro phenomena in pain-based learning, without delving into the underlying neural architecture that such probabilistic models aim to account for [61]. Although increasing evidence supports the role of such realist applications in perception, [39,42,62], including in pain [63,64], these are yet to mature into adequate models of complex learning scenarios. Initial investigations comparing models of learning in pain, including generic Bayesian models [65], suggest that there is work to be done to outperform temporal difference models in computational neuroscience paradigms [66]. Consequently, some have argued that the Bayesian Brain should be treated as an instrumental theory in lieu of more developed mechanistic explanations [67]. An important question for the future application of probabilistic models relates to the nature of our experimental paradigms in pain. Using a model, designed to reflect an active learner who minimises uncertainty over time, may demand an alteration in traditional stimulus-response protocols.

Associative learning theories would be considered incomplete without accounting for value, reward or utility in relation to optimal behaviour. Bayesian generalisations of the Resorla-Wagner model, embodied in the Kalman filter, assumes that the target of learning is the problem of predicting immediate reward [68]. However, full active inference accounts aim to replace the notions of reward, value or utility, by subsuming them all within the generative model [13,69]. Whether these concepts can therefore be considered redundant, while still accounting for the complexities of learning in pain and pleasure, is yet to be determined.

Conclusion

We have presented a broad overview of Bayesian models of learning in pain. From this view, the experience of pain involves the continual prediction of the consequences of action in relation to bodily threat. As such, learning in pain is both predictive and active. Although there still exist many challenges to the full implementation of such probabilistic accounts, we propose that at present, Bayesian derivatives (such as predictive processing and active inference) can provide important considerations for researchers and clinicians alike.

Acknowledgements

The authors would like to extend there gratitude and thanks to Dr Max Jones for his insightful discussion in the critique and development of this review.
This research did not receive any specific grant from funding agencies in the public, commercial, or non-for-profit sectors.

References

1. Sutton RS, Barto AG: *Reinforcement learning: An introduction*. MIT Press; 1998.
2. McNamara JM, Houston AI, Collins EJ: **Optimality models in behavioral biology**. *SIAM Rev* 2001, **43**:413–466.
3. Houston AI, Higginson AD, McNamara JM: **Optimal foraging for multiple nutrients in an unpredictable environment**. *Ecol Lett* 2011, **14**:1101–1107.
4. Vlaeyen JWS, Linton SJ: **Fear-avoidance model of chronic musculoskeletal pain: 12 years on**. *Pain* 2012, **153**:1144–1147.
5. Karos K, Williams AC de C, Meulders A, Vlaeyan J: **Pain as a threat to the social self: a motivational account**. *Pain* 2018, **In Press**.
6. Fordyce WE: **Behavioural science and chronic pain**. *Postgrad Med J* 1984, **60**:865–868.
7. Meulders A, Vansteenwegen D, Vlaeyen JWS: **The acquisition of fear of movement-related pain and associative learning: A novel pain-relevant human fear conditioning paradigm**. *Pain* 2011, **152**:2460–2469.
8. Vlaeyen JWS: **Learning to predict and control harmful events: Chronic pain and conditioning**. *Pain* 2015, **156**:S86–S93.
9. O'Doherty JP, Seymour B, Koltzenburg M, Frackowiak RSJ, Wiech K, Friston KJ, Dolan RJ: **Opponent appetitive-aversive neural processes underlie predictive learning of pain relief.** *Nat Neurosci* 2005, **8**:1234–40.
10. O'Doherty JP, Dayan P, Friston K, Critchley H, Dolan RJ: **Temporal difference models and reward-related learning in the human brain**. *Neuron* 2003, **38**:329–337.
11. Seymour B, Daw ND, Roiser JP, Dayan P, Dolan R: **Serotonin Selectively Modulates Reward Value in Human Decision-Making**. *J Neurosci* 2012, **32**:5833–5842.
12. Seymour B, O'Doherty JP, Dayan P, Koltzenburg M, Jones AK, Dolan RJ, Friston KJ, Frackowiak RS: **Temporal difference models describe higher-order learning in humans**. *Nature* 2004, **429**:664–667.
13. Friston KJ, Daunizeau J, Kiebel SJ: **Reinforcement learning or active inference?** *PLoS One* 2009, doi:10.1371/journal.pone.0006421.
14. Dayan P, Berridge KC: **Model-based and model-free Pavlovian reward learning: Revaluation, revision, and revelation**. *Cogn Affect Behav Neurosci* 2014, **14**:473–492.
15. Gershman SJ, Daw ND: **Perception, action and utility - the tangled skein**. *Princ Brain Dyn Glob State Interact* 2011,
16. Daw ND, Niv Y, Dayan P: **Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control**. *Nat Neurosci* 2005, **8**:1704–1711.
17. Crombez G, Eccleston C, Van Damme S, Vlaeyen JWS, Karoly P: **Fear-avoidance model of chronic pain: the next generation.** *Clin J Pain* 2012, **28**:475–83.
18. Wiech K, Tracey I: **Pain, decisions, and actions: a motivational perspective**. *Front Neurosci* 2013, doi:10.3389/fnins.2013.00046.
19. Vlaeyen JWS, Crombez G, Linton SJ: **The fear-avoidance model of pain**. *Pain* 2016, **157**:1588–1589.
20. Gatzounis R, Schrooten MGS, Crombez G, Vlaeyen JWS: **Operant Learning Theory in Pain and Chronic Pain Rehabilitation**. *Curr Pain Headache Rep* 2012, **16**:117–126.
21. Melzack R, Wall P: *The Challenge of pain*. Viking Penguin; 1988.
22. Wall PD: **On the relation of injury to pain. The John J. Bonica Lecture**. *Pain* 1979, **6**:253–264.

23. Clark A: *Surfing uncertainty: Prediction, action, and the embodied mind*. Oxford University Press; 2017.
24. Hohwy J: *The predictive mind*. Oxford University Press; 2013.
25. Körding KP, Wolpert DM: **Bayesian integration in sensorimotor learning**. *Nature* 2004, **427**:244–247.
26. Kruschke JK: **Bayesian approaches to associative learning: From passive to active learning**. *Learn Behav* 2008, **36**:210–226.
27. Tabor A, Thacker MA, Moseley GL, Körding KP: **Pain: A Statistical Account**. *PLoS Comput Biol* 2017, **13**.
28. Vilares I, Kording K: **Bayesian models: the structure of the world, uncertainty, behaviour, and then brain**. *Ann N Y Acad Sci* 2011, **1224**:22–39.
29. Clark A: **Embodied Prediction**. In *Open MIND*. Edited by Metzinger T, Windt JM. MIND Group; 2015:7.
30. Doya K, Ishii S, Pouget A, Rao RPN: *Bayesian Brain: Probabilistic Approaches to Neural Coding*. 2007.
31. Friston K: **The history of the future of the Bayesian brain**. *Neuroimage* 2012, **62**:1230–1233.
32. Griffiths TL, Kemp C, Tenenbaum JB: **Bayesian models of cognition**. In *Cambridge handbook of computational cognitive modeling*. 2007:1–49.
33. Pouget A, Beck JM, Ma WJ, Latham PE: **Probabilistic brains: Knowns and unknowns**. *Nat Neurosci* 2013, **16**:1170–1178.
34. Trommershauser J, Kording KP, Landy MS: *Sensory Cue Integration*. Oxford University Press; 2011.
35. Kording KP, Beierholm U, Ma WJ, Quartz S, Tenenbaum JB, Shams L: **Causal inference in multisensory perception**. *PLoS One* 2007, **2**.
36. Koller D, Friedman N: *Probabilistic graphical models: principles and techniques*. MIT Press; 2009.
37. Chater N, Oaksford M: *The probabilistic mind: Prospects for Bayesian cognitive science*. Oxford University Press; 2008.
38. Clark A: **Whatever next? Predictive brains, situated agents, and the future of cognitive science.** *Behav Brain Sci* 2013, **36**:181–204.
39. Hohwy J: **Priors in perception: Top-down modulation, Bayesian perceptual learning rate, and prediction error minimization**. *Conscious Cogn* 2017, **47**:75–85.
40. Seth AK: **The Cybernetic Bayesian Brain**. In *Open MIND*. . 2015.
41. Clark A: **The many faces of precision (Replies to commentaries on "Whatever next? Neural prediction, situated agents, and the future of cognitive science."** *Front Psychol* 2013, **4**.
42. Haarsma J, Fletcher P, Ziauddeen H, Spencer T, Diederen K: **Precision weighting of cortical unsigned prediction errors is mediated by dopamine and benefits**. *bioRxiv* 2018, doi:10.1101/288936.
43. Adams R a, Stephan KE, Brown HR, Frith CD, Friston KJ: **The computational anatomy of psychosis**. *Front Psychiatry* 2013, **4**:47.
44. Friston KJ, Shiner T, FitzGerald T, Galea JM, Adams R, Brown H, Dolan RJ, Moran R, Stephan KE, Bestmann S: **Dopamine, affordance and active inference**. *PLoS Comput Biol* 2012, doi:10.1371/journal.pcbi.1002327.
45. *Friston K, FitzGerald T, Rigoli F, Schwartenbeck P, O'Doherty J, Pezzulo G: **Active inference and learning**. *Neurosci Biobehav Rev* 2016, **68**:862–879.

An active inference accoun of learning. The authors discuss model-based and model free learning in the context of belief. Here, the value function is absorbed into a single functional belief, centred around uncertainty.

46. Friston K, Mattout J, Kilner J: **Action understanding and active inference**. *Biol Cybern* 2011, doi:10.1007/s00422-011-0424-z.
47. Seth AK, Friston KJ: **Active interoceptive inference and the emotional brain**.

Philos Trans R Soc Lond B Biol Sci 2016, **371**.
48. Seth AK: **Interoceptive inference, emotion, and the embodied self**. *Trends Cogn Sci* 2013, **17**:565–573.
49. *Owens AP, Allen M, Ondobaka S, Friston KJ: **Interoceptive inference: From computational neuroscience to clinic**. *Neurosci Biobehav Rev* 2018, **90**:174–183.
The authors consider predictive processing and active inference as potential frameworks to describe the role of interoceptive inference in health and disease. Building upon the EPIC model put forward by Barrett and Simmons, 2015, they assert that pathology may arise as a consequence of deficits in the attenuation of ascending prediction error.

50. Pezzulo G: **Why do you fear the bogeyman? An embodied predictive coding model of perceptual inference**. *Cogn Affect Behav Neurosci* 2014, **14**:902–11.
51. Pezzulo G, Rigoli F, Friston K: **Active inference, homeostatic regulation and adaptive behavioural control**. *Prog Neurobiol* 2015, **134**:17–35.
52. Barrett LF, Simmons WK: **Interoceptive predictions in the brain**. *Nat Rev Neurosci* 2015, doi:10.1038/nrn3950.
53. Edwards MJ, Adams RA, Brown H, Pareés I, Friston KJ: **A Bayesian account of “hysteria.”** *Brain* 2012, **135**:3495–3512.
54. Brown RJ: **Medically unexplained symptoms: a new model**. *Psychiatry* 2006, doi:10.1383/psyt.2006.5.2.43.
55. Van den Bergh O, Witthöft M, Petersen S, Brown RJ: **Symptoms and the body: Taking the inferential leap**. *Neurosci Biobehav Rev* 2017, doi:10.1016/j.neubiorev.2017.01.015.
56. *Clark: **Busting out: predictive brains, embodied minds, and the puzzle of the evidentiary veil**. *Nous* 2016, doi:doi: 10.1111/nous.12140.
An accessible overview of the predictive brain literature, with a poignant attempt to extend the scope from a neurocentric mind to an extended mind. The brain from this account is just part of the embodied architecture that predicts the consequences of action.

57. Friston K: **The free-energy principle: A unified brain theory?** *Nat Rev Neurosci* 2010, doi:10.1038/nrn2787.
58. Bowers JS, Davis CJ: **Bayesian just-so stories in psychology and neuroscience**. *Psychol Bull* 2012, doi:10.1037/a0026450.
59. *Klein C: **What do predictive coders want?** *Synthese* 2016, doi:10.1007/s11229-016-1250-6.
A pessimistic overview of Predictive Coding/Procesing as unifying theories of mind, critiquing the ability of such accounts to accommodate motivation. The use of the dark room problem is central. Klein reconciles the application by providing insights into how such accounts can be enhanced through embodied frameworks.

60. Jones M, Love BC: **Bayesian fundamentalism or enlightenment? on the explanatory status and theoretical contributions of bayesian models of cognition**. *Behav Brain Sci* 2011, **34**:169–188.
61. Rasmussen D, Eliasmith C: **God, the devil, and the details: Fleshing out the predictive processing framework**. *Behav Brain Sci* 2013, **36**:223–224.
62. Gordon N, Koenig-Robert R, Tsuchiya N, Van Boxtel JJA, Hohwy J: **Neural markers of predictive coding under perceptual uncertainty revealed with hierarchical frequency tagging**. *Elife* 2017, **6**.
63. Buchel C, Geuter S, Sprenger C, Eippert F: **Placebo analgesia: a predictive coding perspective**. *Neuron* 2014, **81**:1223–1239.
64. Anchisi D, Zanon M: **A Bayesian perspective on sensory and cognitive integration in pain perception and placebo analgesia**. *PLoS One* 2015, **10**:1–20.
65. Mathys C: **A Bayesian foundation for individual learning under uncertainty**. *Front Hum Neurosci* 2011, doi:10.3389/fnhum.2011.00039.

66. Zhang S, Mano H, Lee M, Yoshida W, Kawato M, Robbins T, Seymour B: **The control of pain by active relief learning**. *Elife* 2018, **7**:e31949.
67. Colombo M, Series P: **Bayes in the brain—on Bayesian modelling in neuroscience**. *Br J Philos Sci* 2012, **63**:697–723.
68. *Kruschke JK: **Bayesian approaches to associative learning: From passive to active learning**. *Learn Behav* 2008, doi:10.3758/LB.36.3.210.

A core text expressing the differences and similarities of traditional associative learning models in relation to Bayesian approaches. A demonstration of how uncertainty and active learning are accommodated by a Bayesian model, accounted for backward blocking phenomena.

69. Friston KJ: **What is optimal about motor control?** *Neuron* 2011, **72**:488–498.
70. Friston KJ, Daunizeau J, Kiebel SJ: **Reinforcement learning or active inference?** *PLoS One* 2009, **4**.