



Hierarchical models of pain: Inference, information-seeking, and adaptive control.

Ben Seymour^{a,b,*}, Flavia Mancini^a

^a Computational and Biological Learning Lab, Department of Engineering, University of Cambridge, United Kingdom

^b Center for Information and Neural Networks, National Institute of Information and Communications Technology, Osaka, Japan

ARTICLE INFO

Keywords:

Pain
Nociception
Information theory
Reinforcement learning
Optimal control
Predictive coding
Epistemic value
Free energy principle
Endogenous modulation

ABSTRACT

Computational models of pain consider how the brain processes nociceptive information and allow mapping neural circuits and networks to cognition and behaviour. To date, they have generally have assumed two largely independent processes: perceptual inference, typically modelled as an approximate Bayesian process, and action control, typically modelled as a reinforcement learning process. However, inference and control are intertwined in complex ways, challenging the clarity of this distinction. Here, we consider how they may comprise a parallel hierarchical architecture that combines inference, information-seeking, and adaptive value-based control. This sheds light on the complex neural architecture of the pain system, and takes us closer to understanding from where pain 'arises' in the brain.

1. Introduction

One of the most fundamental questions in pain neuroscience is identifying what brain activity is necessary and sufficient for the conscious experience of pain. A common conceptual approach is to try and deconstruct different components of pain behaviour and experience, and then re-piece them together to understand how pain emerges from their coordinated operation (Melzack, 1968). But this leads to the problem of identifying precisely what these individual components should be, and then understanding how they functionally fit together. In contrast, computational approaches define the specific computations that different neural populations implement, and then determine how they are integrated to generate pain experience and behaviour. Mathematical accounts have focused on two central processes: perceptual inference (i.e. identifying the sensory features and causes of a pain state) and action control (i.e. evoking responses or actions appropriate to the event). Corresponding models, in particular Bayesian inference and reinforcement learning respectively, have had some success in explaining behaviour and corresponding neural activity across a range of appropriately-designed experimental paradigms (Seymour and Dolan, 2013; Büchel et al., 2014; Wiech, 2016). But this leads to the question of how they interact, and to what extent they are really independent processes, especially as the subjective experience of pain seems to seamlessly integrate both perceptual and motivational features.

In this review, we first outline the conceptual basis for computational models of pain. We review theories of perceptual inference (Bayesian and predictive coding models) and action control (reinforcement learning). We then discuss how these two processes interact, and what this means for our understanding of pain perception and behaviour.

2. The computational approach to pain

Computational models aim to provide a *mechanistic* understanding of how activity in the brain yields behaviour. Ultimately, this views brain activity in terms of information processing: what quantities are represented by the brain, and how the brain computes information to generate behaviour in the organism. Marr provided a general framework for such a computational approach (Marr et al., 1979), outlining 3 basic levels of understanding that have been enormously influential in neuroscience. Marr's three levels of understanding brain systems comprise: (1) a computational level - understanding formally what problems the system is trying to solve; (2) an algorithmic level - what solutions the system adopts at a mathematical level; and (3) an implementational level - how these mathematical operations are achieved by biological and neuronal 'hardware'. Any complete understanding of the pain system should span all 3 levels (Fig. 1).

Marr's conceptualisation of different levels illustrates that there are two different ways in which computational models can be applied to a biological system such as pain (Dayan and L.F. Abbott, 2001; Kaplan, 2011). At the computational level, for instance, one can pro-

* Corresponding author.

E-mail addresses: ben.seymour@ndcn.ox.ac.uk (B. Seymour), flavia.mancini@eng.cam.ac.uk (F. Mancini).

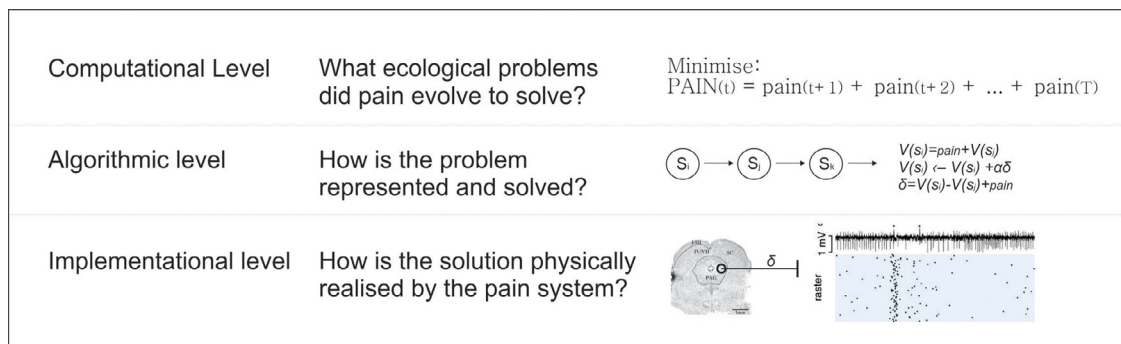


Fig. 1. Computational neuroscience and Marr's 3 levels. Marr proposed a foundational approach for understanding neural systems based on computational, algorithmic and implementational levels. In the example on the right, for simple pain conditioning, the overall computational goal is defined as the minimisation of pain. Although there are several algorithms which can achieve this, a biological plausible one comes from reinforcement learning (temporal difference learning), and involves error-based learning an aversive value metric for discrete states. Neurophysiological evidence shows that PAG neurons actually encode a corresponding error signal (Johansen et al., 2010), illustrating overall how the brain can achieve the required behaviour.

pose models that achieve the overall aim of the system and replicate observable behaviour, but don't necessarily describe the underlying algorithmic process itself: in this sense, these can be thought of as *descriptive* models, because they capture the problem the brain is trying to solve. However, one can also propose models that aim to emulate both observable behaviour and the actual computations that the brain directly and physically implements: these can be thought of as *mechanistic* models, because they solve the problem in the same way that the brain does. Clearly, the goal of computational neuroscience is to determine the most accurate mechanistic models, because this can lead to an understanding of how the brain implements computations in terms of physiological processes at the level of molecules, neurons and circuits. In this way, the model provides a fundamental level of understanding that describes the operation of the system, that transcends mere descriptions of physiological interactions at various neuroanatomical scales.

3. Sensory processing of pain as statistical inference

The idea that sensory processing of pain signals in the brain might comprise some sort of statistical inference about the cause of incoming nociceptive information is a longstanding idea, following parallels with other sensory domains such as vision and audition (Wager, 2005; Watson et al., 2006; Brown et al., 2008, 2008; Seymour and Dolan, 2013; Yoshida et al., 2013; Moutoussis, Fearon, et al., 2014; Anchisi et al., 2015; Wiech, 2016; Tabor et al., 2017; Ongaro et al., 2019; Hoskin et al., 2019). Inference models, especially Bayesian models, have become popular in the context of the broader 'Bayesian brain hypothesis', and reflect the fact that sensory percepts are typically biased by sources of predictive information in an approximately Bayesian optimal way, such that the resulting percept reflects the integration of this information (Knill et al., 2004; Friston, 2012). This information can come from predictions generated by preceding (prior) information or cues (e.g. the sight of a dentist drill going into your mouth), or multisensory integration (e.g. seeing blood coming from a fresh wound).

In the case of pain, the argument in favour of statistical inference models has primarily come from the observation that pain intensity judgements are biased in the direction of expectations, which is a defining feature of placebo and nocebo effects (Atlas, Bolger, et al., 2010; Colloca et al., 2006; Brown et al., 2008; Wiech, Vandekerckhove, et al., 2014). From a Bayesian perspective, inference involves integrating an expected distribution of intensity (the 'prior') with noisy incoming sensory information (the 'likelihood') in a statistically optimal way, and this yields the inferred probability distribution of the actual intensity (the 'posterior'). Ultimately, experienced pain should reflect some function of the statistics of this posterior distribution (Tabor et al., 2017).

Evidence for Bayesian inference in pain processing was first formally investigated by Yoshida et al. (2013), who looked at the influence of uncertainty of the prior distribution of intensity on phasic thermal pain. Although this clearly supported the notion that pain involves some sort of inference, this and subsequent studies also indicated that the nature of this inference may be quite complex. For instance, uncertainty appears to enhance pain *per se*, over and above its role in inference, suggesting that the brain adds a 'cost' to uncertainty as a post-processing step to inference (Yoshida et al., 2013; Anchisi et al., 2015). Furthermore, the prior precision can be resistant to updating, in the case where explicit incorrect predictions are continually refuted by experience (Jepma, 2018). Further still, rather than a smooth integration of priors and likelihoods, priors can be immediately down-weighted if there is a large discrepancy with the sensory evidence, suggesting a thresholding process that detects sudden changes in the environment (unexpected uncertainty (Yu and Dayan, 2005)), to re-evaluate the world if one's internal model is wrong (Hird et al., 2018). Together, these issues illustrate the problem if know what exactly the brain is optimising, with what cost function, and based on what information (Wald, 1947).

Previous studies have focused on uncertainty in the prior distribution, but the importance of uncertainty in the likelihood (i.e. the noisiness of the nociceptive input) has not been well studied because it is difficult to manipulate experimentally (unlike Gabor gratings in vision, for example). Furthermore, other aspects of pain have not been formally investigated, notably modulation by multisensory or motor efference information. In the case of multisensory modulation, phenomena including touch-induced analgesia (Mancini, Beaumont, et al., 2015), modulation by colour information (Moseley et al., 2007) and the painful rubber hand illusion (Ehrsson et al., 2007) illustrate how pain is modulated by other sensory modalities, but there have been no systematic computationally motivated studies investigating the influence of concomitant sensory stimuli - something that is likely to be important in naturally occurring nociceptive stimuli (Haggard et al., 2013). Similarly, the role of motor efference information is likely to be important: for instance, self-induced stimulation reduces perceived pain, but it is unclear if this relates to timing uncertainty or cognitive controllability and self-efficacy (Braid et al., 2006; Litt, 1988; Mohr et al., 2005; Helmchen et al., 2006; Wiech, Kalisch, et al., 2006; Salomons et al., 2007; Bräscher et al., 2016; Borhani et al., 2017).

These issues raise a deeper question as to exactly *what* is being inferred during pain perception. Classical sensory inference tasks are exteroceptive, in which the goal of perception is the identity or an observable and thus externally verifiable property of an external object. In the case of nociception, it makes sense for nociceptive input to be used, at least in part in an exteroceptive manner, because this can resolve external information that may have a causal influence on future events (e.g. iden-

tifying an object on the basis that it caused pain, such as distinguishing a paper clip from a drawing pin by touch) (Price et al., 2003). However, pain also embodies an interoceptive sense concerning inference about internal bodily states: these are intrinsically linked to innate motivational drives supporting homeostasis (Craig, 2009; Farkas, 2013), and not externally observable or verifiable by others. Importantly, the pain system has distinct afferent nociceptive pathways from deep and superficial laminae of the dorsal horn of the spinal cord which may relate to dissociable exteroceptive and interoceptive sensing (sending projections to lateral and medial thalamic nuclei, respectively) (Dostrovsky et al., 2013). Although it remains unclear how these parallel pathways contribute to the overall subjective perception of pain, this suggests that there may be distinct inference processes underlying interoceptive and exteroceptive pathways.

3.1. Predictive coding

Bayesian models are typically applied to perception at a descriptive level. Whereas perception may be approximately Bayes optimal in some cases, it doesn't necessarily mean that the brain is implementing Bayesian inference directly (Colombo et al., 2012). One reason for questioning the plausibility of Bayesian models at an algorithmic level is that representing full probability distributions across a neural population is unrealistic unless the distribution is simple, which is rarely likely to be the case. This has led to algorithmic approaches to perceptual inference that approximate inference, such as Predictive Coding (Rao et al., 1999; Friston, 2005; Aitchison et al., 2017). In predictive coding models, inference occurs on a hierarchy in which higher level features are inferred on the basis of lower level features. A particular level encodes a prior expectation of what its input will be, and sends this information to the level below. This lower level compares this prediction with the observed features and computes a prediction error; the prediction error is sent back to the higher level and can be used to update the prior expectations to minimise surprise (Fig 2). The amount of belief updating is determined by the relative precision of prior and incoming information, allowing information to be weighted by its uncertainty (Fig 2) (Friston, 2008). For completeness, it is worth noting that there are notions of predictive coding that are not fully *internalistic*, incorporating enactive, embodied and extended cognition (Allen et al., 2018; Ramstead et al., 2019; Bruineberg et al., 2018).

In the case of pain, predictive coding may offer an account of the nature of both exteroceptive and interoceptive pain processing in different cortical hierarchies (Büchel et al., 2014; Fardo et al., 2017; Owens et al., 2018; Seth, 2013; Allen et al., 2018). In the context of interoceptive processing, it is known that posterior, granular regions of insula cortex receive afferent input arising from superficial laminae of the dorsal horn, projecting via medial thalamic nuclei (including VMPo) to the insula, where other visceral and interoceptive afferents also converge (Craig, 2002; Evrard, 2019). Within the insula, there appears to be a hierarchical remapping along a posterior to anterior gradient, with more complex, high-level representations in the anterior insula (Evrard et al., 2014), and this could plausibly serve as a substrate for an interoceptive predictive coding hierarchy. Neuroimaging studies provide provisional support for this hypothesis: for instance, one would predict that neural activity in layers of the hierarchy that primarily code prediction errors, as opposed to features themselves, would evoke less activity for expected than unexpected pain (Geuter, 2017).

It has also been suggested that the predictive coding hierarchy may extend to the brainstem and even the spinal dorsal horn, with potential insight into the nature of ascending and descending architecture of pain pathways (Büchel et al., 2014). Some evidence for this comes from evidence of precision coding (i.e. uncertainty) in the periaqueductal grey (PAG) during inference tasks (Grahel et al., 2018). The PAG is a critical hub for pain processing, receiving major descending projections from sub-regions of the insula cortex, as well as ascending input from the dorsal horn, to which it also reciprocally projects via the rostroventral

medulla (Basbaum et al., 1984; Heinricher et al., 2009; Evrard et al., 2012).

In the context of exteroceptive processing, it is known that the primary somatosensory cortex (S1) and secondary somatosensory cortex (S2) receive projections that arise from lamina I and deeper layers of the dorsal horn via more lateral thalamic nuclei, where they converge with other exteroceptive somatosensory input (Apkarian et al., 1989; Dostrovsky et al., 2013; Stevens et al., 1993). These cortical projections are topographically organised (Mancini, Haggard, et al., 2012), presumably supporting fine spatial encoding and discrimination of nociceptive inputs (Mancini, Bauleo, et al., 2014). S1 and S2 are mutually interconnected (Pandya et al., 1969); they both send and receive projections to/from subregions of the posterior parietal cortex, including the superior parietal lobule (areas 5 and 7), the inferior parietal cortex, the insula, and the motor and premotor cortex (Darian-Smith et al., 1993; Flaherty et al., 1991; Pons et al., 1986). The primate S1 also sends somatotopically-arranged projections onto the putamen (Flaherty et al., 1991, 1995; Robbe, 2018). Very few studies have investigated whether this neural pathway supports predictive coding of pain signals and the little evidence available is controversial. In support of the predictive coding view, it has recently been shown using MEG that spatial pain processing within this hierarchy is modulated by expectations and their violations in a manner consistent with a predictive coding model (Fardo et al., 2017). However, a recent fMRI study suggests that neural activity in S1 and S2 primarily relate to sensory encoding rather than predictive coding of nociceptive signals (Geuter, 2017), although it did report evidence of predictive coding in the anterior insula.

In summary, the notion that pain involves a sensory processing stream that performs some sort of statistical inference has preliminary support but needs further research. It is plausible that statistical inference occurs in pathways that are at least partly parallel and partly distinct along spino-thalamo-cortical hierarchies for interoceptive and exteroceptive processing. At the same time, the cortical neuroanatomy is clearly more complex than can be fully explained by current computational models (Vierck, 2013). Furthermore, it is also clear that the subjective experience of pain cannot be solely explained by statistical inference of sensory signals, which is further evident in experiments that probe the motivational basis of pain.

4. Motivational processing of pain as optimal control

Pain invariably leads to some sort of action, ranging from simple motor reflexes and physiological responses, to conscious 'deliberative' decisions. Such actions should be appropriate to the pain experienced, either preparing for, reducing, or completely avoiding it if possible (Fields, 2018). Optimal actions are usually dependant on learning, which has been well studied in terms of Pavlovian (Classical) and Instrumental (Operant) learning (Mackintosh, 1983). This has uncovered much about the structure of motivation and learning systems across species, exposing how observed behaviour derives from multiple underlying interacting control processes.

First, pain induces a set of innate, evolutionarily 'hard-wired' responses: some of these are specific to pain (e.g. limb withdrawal), some generalise to aversive stimuli of any cause (e.g. withdrawal, autonomic arousal), and many are remarkably sophisticated, context-sensitive behavioural repertoires (Bolles, 1971).

Second, this innate response set can be hijacked through Pavlovian associative learning, such that a cue that reliably predicts pain will elicit a response similar to the innate response on presentation of the pain itself. This is not implemented by a single learning system, but by multiple, partially independent learning processes (e.g. cerebellar learning of precisely timed motor limb withdrawal or eye-blink responses (Zhang et al., 2016; Ernst et al., 2019), and amygdalar learning of autonomic responses, freezing and other generalised aversive responses (Herry et al., 2014; Janak et al., 2015)).




Fig. 2. Simplified model of predictive coding and reinforcement learning. On the left, the predictive coding model sees predictions (the prior) transmitted to the level below, in anticipation of an incoming signal. This is compared to the signal (the likelihood) and prediction error computed and sent back to the level above to modify the expectation (though learning). The learning process may be modified by a precision signal (not shown) in keeping with a (Bayesian) predictive coding model (Friston, 2008; Aitchison et al., 2017). On the right, the reinforcement learning (RL) model considers three classes of action at different levels. On the bottom, incoming nociceptive signals reflexively lead straight to an innate response (although they do not require it, they may be modulated from 'higher-order' control centers). The intermediate 'model-free' level involves value learning of states and actions, where the values simply represent the overall goodness or badness of a state or action. As in predictive coding, these are learned using a prediction error, and well described by simple RL models such as Temporal Difference learning (Sutton et al., 1998; Dayan, L. Abbott, et al., 2003). At the highest 'model-based' level, action control involves an internal representation of state and action space, which can be flexibly used for learning, inference and control. This can be achieved via a broad array of algorithms, including Bayesian and Dynamic Programming approaches. In the figure, we also sketch how the two might fit together, with inferred pain signals being transferred from the predictive coding hierarchy to the RL hierarchy, although as we describe below, this is an over-simplification.

Third, animals can learn to actively escape from and avoid pain through directly associating pain with specific preceding actions. This kind of learning, termed instrumental learning, can involve a sophisticated interplay between Pavlovian learning and reward learning processes, involving inhibition of actions that predict pain and selection of other responses that lead to relief, ideally before pain occurs (Mowrer, 1960; Mackintosh, 1983). Although such escape and avoidance actions can often be achieved purely by Pavlovian responses (Bolles, 1972; Mackintosh, 1983), instrumental learning allows a much greater diversity of action types to be learned, i.e. actions that wouldn't normally be within the Pavlovian / innate repertoire. In the basic instrumental case, termed 'model-free' learning, actions are emitted simply according to their learned value (i.e. on a scale from good to bad), but don't evoke any representation of the actual outcome (Dayan and Daw, 2008). In this way, such instrumental actions are sometimes called stimulus-response 'habits' (Dickinson and Balleine, 2002).

Fourth, it is also clear that learning and action selection, especially in humans, can be enormously sophisticated, involving a consciously accessible representation of exactly how events or actions lead specifically to pain and other outcomes. This sort of internal 'model-based' learning reflects what is normally considered to be cognitive or deliberative action selection, and allows great flexibility for mental planning, counter-factual thinking, episodic learning, contingency knowl-

edge, and observational and instructed learning (Atlas, Doll, et al., 2016; Koban et al., 2019). Because of its potential complexity, cognitive decision-making is the least understood, but clearly works alongside the other action systems: usually in concert, but occasionally in conflict with them (Carter et al., 2006; Dayan, Niv, et al., 2006).

4.1. Reinforcement learning

Given the complexity of the environment, most actions need to be learned or inferred from experience. In this context, pain acts as a teaching signal to optimise future behaviour in terms of harm limitation. But this is a difficult problem, because optimising action incorporates the fact that actions should be employed as early as they can reliably avoid or minimise pain. And since most interactions with the world involve complex sequences of individual actions, the brain needs to determine which actions in any sequence are the key actions that lead to avoidance, and which are inconsequential. This is called the 'credit assignment problem' and falls within a domain of control theory called Reinforcement Learning (RL) (Bellman, 2013). RL captures the nature of the learning and action selection through trial-and-error interaction with the environment, and has led to a number of algorithmic solutions of varying complexity and potential biological plausibility (Sutton et al., 1998).

A particular class of algorithm, called Temporal Difference learning (TD) and its variants, has attracted substantial attention in neuroscience, because it offers an intuitive account of many well-studied learning phenomena in Pavlovian and instrumental learning (for both rewards and punishments) (Sutton et al., 1981). It involves learning a value term for any cue or action, computing a value prediction error as one moves through time (the difference between what happened and what was expected), and using the prediction error to update preceding values if their effective predictions are refuted by subsequent experience to a greater or lesser extent. Critically, it has been shown that phasic dopaminergic activity in the ventral tegmental area and substantia nigra encode reward prediction error signals (Schultz et al., 1997); this finding led to the notion that these models directly implement RL algorithms, offering a mechanistic explanation of behaviour (Schultz, 2016).

In the case of pain, there is reasonable evidence that the brain implements TD-like algorithms for basic Pavlovian value learning of phasic pain, as well as relief of tonic pain. In particular, ventral putamen responses correlate with a higher order pain prediction error, illustrating how prediction errors can be passed in chains of predictive cues before pain occurs (implying that they are not just tied to pain outcomes) (Seymour et al., 2004, 2005). Instrumental avoidance is more complex because it necessarily involves the interaction between state and action learning processes. Integrated TD-like models (such as the actor-critic model) provide a good account at a theoretical level (Maia, 2010; Moutoussis, Bentall, et al., 2008); within these models, a basic TD learning process concurs with the instrumental value learning component of both avoidance and escape learning (Seymour et al., 2012; Roy et al., 2014).

Some Pavlovian phenomena, such as latent inhibition, are difficult to explain by simple RL algorithms, and almost certainly require an internal, generative model of the environment (Gershman and Niv, 2012). This is also the case for much of the more deliberative and cognitive aspects of pain avoidance (Tolman, 1948; Gillan et al., 2016). It is often difficult to determine exactly how the brain learns and maintains complex internal models, but Bayesian models - in which one optimally infers state and action transition probabilities from experience - typically provide an intuitive approach to understand behaviour at a descriptive level, at the very least. Conceptually, an internal model allows the brain to predict and evaluate sequences of actions and events. In this sense, action planning inverts this sequence given a desired outcome, to try to determine what action is necessary now to optimise the chances of a given state, such as successful avoidance of pain (Botvinick and Toussaint, 2012).

It remains the case, however, that on the surface Bayesian RL algorithms often don't appear to capture choice behaviour in simple escape and avoidance experiments, illustrating that it can be very hard to know exactly what the brain is optimising, and that one needs to be cautious when looking at a proposed circuit in isolation given that there are multiple mechanisms that the brain can utilise to control behaviour. This leads to the notion that emitted behaviour derives from the interplay between a complex hierarchy of learning processes (Daw et al., 2005; Lee et al., 2014). Behavioural evidence supports this integration process in the case of pain (Wang et al., 2018).

More broadly, the hierarchy of controllers that spans simple innate responses to complex cognitive decision-making illustrates the trade-off between speed and computational simplicity - which is critical for balancing rapid defensive responding, accuracy and computational sophistication (O'Doherty et al., 2015). It seems likely that the structure of pain motivational systems reflects the evolutionary addition of successively more sophisticated components.

4.2. Information-seeking

The dynamics of avoidance also raise a further issue that gets to a fundamental problem in learning and motivation. If an action leads to successful avoidance (i.e. no externally discernible outcome), then how

long should the action be maintained over time? If avoidance is maintained indefinitely, then one will never know if the action is forever necessary, or if the avoided outcome has otherwise subsided (Denrell et al., 2001). This problem illustrates a general problem in RL, called the exploration-exploitation dilemma, and illustrates why many RL algorithms are designed to consider two types of value: that related to the outcome itself, and that related to information. Outcome-driven actions are exploitative, because they relate to the best course of action given what is currently known (Sutton et al., 1998). However, information-driven actions are exploratory, because they are directed at finding out if there are better actions whose current outcome value is partially or fully unknown (Thrun, 1992; Dayan and Sejnowski, 1996; Kakade et al., 2002). Exploration is best studied in the context of rewards, revealing three basic schema employed by the brain that relate to different controllers (Gershman, 2018): (1) a simple, novelty-seeking heuristic that drives actions based on perceptual novelty (Wittmann et al., 2008); (2) 'undirected' exploration, that effectively adds a level of noise to choices, sometimes driving choice away from the currently highest valued (Daw and O'Doherty, 2006); and (3) 'directed' exploration that motivates action based on the calculated information content (related to uncertainty) of an outcome (Wilson et al., 2014).

Exploration is not well studied in the case of pain-motivated action, although there is evidence of basic noisy exploration during learning, representing a type of 'model-free' information-seeking (Seymour, Daw, et al., 2012; Zhang et al., 2018). Importantly, however, uncertainty during learning appears to increase perceived pain (Yoshida et al., 2013; Taylor et al., 2017; Zhang et al., 2018). This implies that exploration may shape the experience of pain itself, over and above driving exploratory actions. Although an explicit link between uncertainty-based modulation of pain perception and exploration has yet to be shown, this finding provides support to the idea that experienced pain is modulated by not just the direct outcome value associated with the stimulus, but also its informational value. Furthermore, the effect of uncertainty on pain may also be modulated by controllability, suggesting that informational value may be enhanced when it is exploitable (Zhang et al., 2018). This leads to a control system in which pain experience is associated to a flexible, adaptive value tuned by the long-run utility of its information for survival, rather than a fixed outcome value estimate of potential harm (Seymour, 2019).

5. How do sensory and motivational processes interact?

According to the theoretical propositions above, perceptual inference and action control represent distinct computational processes, both utilising computational architectures that are at least partially hierarchically organised. This leads to the question of how they interact. Intuitively, the system encoding perceptual inferences about pain could transfer its output to an action control system, and vice versa, at comparable levels of their respective hierarchies: in simple words, this would allow the brain to decide 'what is this?', 'how harmful is it?' and 'what shall I do about it?', at varying levels of computational speed and sophistication.

There is some evidence for this idea. At the lowest level, when nociceptive signals enter the dorsal horn, there are clearly segmental circuits that mediate basic motor reflexive responses, which are emitted before the brain processes nociceptive information (Todd, 2010). At the highest level, perceptual awareness of a pain state should relate to a cognitive, model-based system that yields reportable ('stated') preferences about proposed action. This must be the case even if consciously experienced pain is not always acted upon - a situation that can exist because 'lower' controllers (i.e. model-free habits, Pavlovian and innate responses) can still exert a dominant effect on overall behaviour (which causes the well known distinction between wanting and liking in the case of rewards (Berridge, 2009), or unwanting and disliking in the case of punishment/pain (Seymour and Dolan, 2013)).

Interactions between inference and control processes at intermediate levels of respective processing hierarchies are harder to discern. However, there are a number of plausible candidate mechanisms. Firstly, it is reasonable to propose that PAG-basal ganglia-amygdala-thalamic circuits might subserve a subcortical network in which semi-processed nociceptive information can feed into model-free Pavlovian and instrumental control systems. The core input site into model free motivational circuits is via the PAG and parabrachial nuclei (PBN) in the brainstem; both receive direct projections from a population of superficial dorsal horn cells, distinct from projections to the medial thalamic nuclei (including VMPo) (Gauriau et al., 2002; Craig, 2003; Roeder et al., 2016). PAG and PBN also receive strong projections from the insula cortex (Evrard, 2019). Furthermore, the primate amygdala projects to the insula, S2 and S1 (Amaral et al., 1984). This is crucial, because it provides the opportunity for crosstalk between the perceptual inference and control streams. Thus, from an anatomical perspective, the motivational circuits for model-free Pavlovian and instrumental learning receive input both from afferent pathways that are distinct to that involved in perceptual inference, and also from cortical regions that are likely to compute interoceptive and exteroceptive inference.

Other candidates for integration between sensory and motivational pain processing pathways are the anterior insula, putamen, anterior cingulate cortex (ACC), dorsolateral prefrontal cortex (dlPFC), posterior parietal cortex and cerebellum (Fig 3). The anterior insula is highly connected to the posterior insula (which in turns receive projections from sensory pain processing streams, i.e. S1, S2, and parietal cortex), and many regions that form part of broader motivational networks (Cauda et al., 2011); this places the anterior insula in the ideal setting to integrate preprocessed nociceptive information with motivational information (encoded in value-based learning systems); indeed a recent human, neuroimaging study provides support for this hypothesis (Geuter, 2017).

In the basal ganglia, neural populations in the primate putamen receive somatotopically-organised projections from S1 and S2 (Flaherty et al., 1991, 1995; Robbe, 2018). There is also evidence that the ventral putamen encodes prediction errors in a range of both appetitive and aversive learning tasks (Seymour et al., 2004, 2005; Delgado et al., 2008); this does not necessarily imply an interaction with the sensorimotor neural populations of the putamen, but it makes it highly likely.

Another key integrative region is the ACC, a well-known central hub in pain processingreceiving projections from both the insula and the posterior parietal cortex. The ACC is activated during tasks that require to discriminate the intensity and spatial location of nociceptive stimuli (Oshiro et al., 2009). Distinct regions of the ACC are also core nodes in motivational circuits, potentially allowing the ACC to compute the informational value of pain states, and so tune the overall perception of pain according to its ultimate function as a control signal for optimal action (Seymour, 2019).

A region highly connected to the ACC, the dlPFC, is also a candidate for sensory-motivational integration, being engaged in pain discrimination tasks (Bornhövd et al., 2002; Oshiro et al., 2009; Lobanov, 2013), endogenous pain suppression (Lorenz et al., 2003; Wager, Rilling, et al., 2004), encoding value-based representations and top-down action selection (Botvinick and Braver, 2015; Kounieher et al., 2009). Interestingly, there is increasing evidence of dlPFC changes in some chronic pain disorders, which reverse following pain resolution (Seminowicz et al., 2017).

Prefrontal regions, including the dlPFC, are densely connected to the posterior parietal cortex. Amongst other things, the posterior parietal cortex is involved in attention and orientation to the sensory features of nociceptive stimuli (Oshiro et al., 2009; Lobanov, 2013). Neurons in the primate ventral intraparietal cortex can elicit nocifensive responses (Kaas et al., 2016), suggesting that this region is important for integrating nociceptive and motor information for bodily defence. In humans, intraparietal regions appear to also encode RL signals (O'Doherty et al., 2015); in particular, neural activity in the intraparietal cortex correlates




Fig. 3. Simplified model of functional anatomy of pain system. The figure shows major brain regions underlying pain perceptual and action control. The red regions show major pathways involved in exteroceptive and interoceptive perceptual inference, arising from deep and superficial layers of dorsal horn. Distinct ascending pathways also project to the PAG and PBN, which connect widely to brain regions involved in motivational processing (light blue), including amygdala, striatal and basal ganglia nuclei, which access more sophisticated representations of value in lateral orbital (state-based) and ventromedial (action-based) prefrontal cortices. One possible role of these circuits is to provide input that allows the ACC to compute informational value, which tunes the overall perception of pain according to its ultimate function as a control signal for optimal action. Other candidates for integrating perceptual and motivational information are shown in purple, whereas area known to be important for the control of pain behaviour are shown in yellow. Higher level representations of pain may be stored in hippocampal-prefrontal networks (navy), allow conscious episodic memory, conceptual and configural understanding of pain, planning, and communication. aINS: anterior insula; ACC: anterior cingulate cortex; aMCC: anterior middle cingulate cortex; Amy: amygdala; BNST: bed nucleus of the stria terminalis; Cb: cerebellum; dlPFC: dorsolateral prefrontal cortex; DS: dorsal striatum; Hi: hippocampus; HN: habenula; Hypo: hypothalamus; IOFC: lateral orbitofrontal cortex; NAc: nucleus accumbens; PAG: periaqueductal grey; PBN: parabrachial nuclei; pINS: posterior insula; PPC: posterior parietal cortex; RVM: rostral ventromedial medulla; S1: primary somatosensory cortex; S2: secondary somatosensory cortex; Th: thalamus; vmPFC: ventromedial prefrontal cortex; VTA: ventral tegmental area. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

with state-prediction errors, which are used to underpin the learning of state transition probabilities in model-based RL (Gläscher et al., 2010). This raises the interesting possibility that state-prediction errors might be used as a common currency between pain inference and action control.

Finally, the cerebellum is a likely candidate for integrating perceptual and motivational information. Primary afferents transmit nociceptive signals to the cerebellum via the pontine nuclei and the inferior olive (Snyder et al., 1978); in humans, acute and pathological pain activates vermal lobules IV/V and bilateral hemispheric lobule VI (Moulton et al., 2010). These cerebellar regions also receive descending afferent input from several regions involved in sensory and motivational processing (S1, superior parietal lobule and PAG via the inferior olive; amygdala, dlPFC, intraparietal cortex via the pontine nuclei); the cerebellum sends ascending output to the thalamus and hypothalamus (Moulton et al., 2010), which is involved in the regulation of pain behaviour. Although it is not clear how the cerebellum encodes nociceptive information, it seems to have a modulatory effect on pain detection (Siegel et al., 1974) and nocifensive motor behaviour, such as withdrawal and freezing in rodents (Cerminara et al., 2009; Dimitrova et al.,

2003). There is also evidence of cerebellar involvement during the anticipation of painful stimuli (Ploghaus et al., 1999). The cerebellar hemispheric lobules I, IV-VI and the tonsils are also involved in Pavlovian pain conditioning (Lange et al., 2015; Zhang et al., 2016; Frontera et al., 2020; Ernst et al., 2019). Altogether, this evidence suggests the cerebellum might perform (some kind of) predictions about pain that can be relayed to sensory, motivational and motor control pathways.

From a computational perspective, a more subtle issue then relates to exactly what quantities might be shared between perceptual and action/motivational systems. For instance, what information is necessary or sufficient for effective communication to be shared between perceptual and motivational processing streams? Clearly both predictive coding and reinforcement learning algorithms use prediction errors, which raises the question as to whether they encode prediction errors in the *same* way. On the surface, this seems plausible: for instance in a conditioning paradigm whereby a cue predicts pain, the prediction error guides both perception and response acquisition. However, when one looks beyond simple paradigms, the equivalence between a sensory prediction error and a value prediction error becomes problematic.

First, a key aspect of RL is the fact that prediction errors occur at any point antecedent of pain in a sequence or predictors, rather than only at the time of the pain itself. That is, RL predictions for pain will occur when different cues disagree about the prediction of pain, a long time before the pain itself occurs. Indeed this distinguishes early models of Pavlovian conditioning, such as the Rescorla-Wagner model, with RL models, since the former cannot accommodate higher-order conditioning phenomena (Sutton et al., 1981; Seymour et al., 2004). In principle, however, perceptual inference can also be predictive, but this aspect of perceptual inference has not been thoroughly investigated in the case of pain.

Second, a core feature of model-free value learning is that it generalises across punishments of different types. This is manifest in a paradigm termed 'transreinforcer blocking', in which a cue fails to acquire a response to one type of punishment (e.g. pain) in the presence of a cue that already predicts a different type of punishment (e.g. aversive noise). Since a perceptual prediction error to noise and pain would be entirely distinct, blocking must necessarily involve a common, shared aversive signal distinct from perceptual predictions (Dickinson and Dearing, 1979).

Third, pain can be conditioned to be to appetitive values. In counter-conditioning, pain acts as a cue for reward, rather than as the outcome in a Pavlovian paradigm, and in this context pain acquires appetitive conditioned responses (Eroféeva, 1921). This means that the onset of a painful stimulus induces a positive reward-like value prediction error, but a pain-like sensory prediction error, again indicating how perceptual inference and action control must necessarily be dissociable.

Despite these necessary differences in intermediate-level sensory and motivational processing, at the higher level there is clearly a strong resonance between the two systems. This is because both draw on a common internal cognitive representation of the nociceptive event. Just how this internal representation of pain is generated and encoded in the brain is not known, and indeed what this representation actually represents in terms of the exteroceptive or interoceptive aspects of the stimulus remains unclear. However, it seems likely that conscious, cognitive appreciation of a painful stimulus incorporates both the sensory and motivational features in a much more integrated manner than lower levels. This reflects a broader issue of how the brain generates perception, motivation and action, addressed by Bayesian, active inference accounts. For instance, the Free Energy Framework proposes that sensation, motivation and action are intrinsically related by their drive to understand the causes of unexpected stimuli and minimise surprise (Friston, 2010; Pezzulo et al., 2015; Parr et al., 2018); these accounts have been applied to pain only rarely (Fardo et al., 2017; Geuter, 2017; Tabor & Burr, 2019) and more research in this field is needed.

5.1. How does the brain give rise to pain experience?

In terms of its phenomenal content, it is possible to suggest three core components of pain that can be at least partially related to specific computational and anatomical substrates. First, pain experience cannot be devolved from its sensory features, such as quality, intensity and spatial location, and this seems to arise from a primarily exteroceptive sensory inferential pathway involving a somatosensory cortex-centred networks (including with parietal, posterior insula and inferior frontal cortices) (Oshiro et al., 2009; Mancini et al., 2016; Fardo et al., 2017). Second, pain is also an interoceptive sensory percept - an 'embodied homeostatic sense' with an inextricable link to defensive action (Craig, 2003), and this maps well to both an insula cortical hierarchy and value-based learning and control circuits (Seth, 2013; Büchel et al., 2014; Allen et al., 2018). Third, pain is finely tuned according to its long-term value as a learning and control signal, incorporating the exploitable information that it conveys: this may be computed across a network of regions centred on the anterior cingulate cortex involved in uncertainty and controllability computations (Seymour, 2019). Importantly, these hierarchical inference and control operations are likely to be integrated across multisensory channels, including interoceptive and exteroceptive nociceptive pathways, alongside other non-nociceptive sensory and action inference pathways (Friston, 2013). Ultimately these underlying circuits are coordinated over a large anatomical area to yield not only effective behaviour, but also the unified and unique subjective qualia of pain.

One can also consider the 'access' component of pain consciousness (Block, 2005), i.e. the ability to cognitively evoke the concept of pain and use it for 'thinking' and planning. Clearly, the brain builds a conceptual model of pain which allows conscious knowledge and imagination, explicit memory, planning, and verbal communication. This internal conscious representation and understanding is distinct from the phenomenal feeling of pain (e.g. thinking about pain is not phenomenologically painful), but allows intelligent defensive behaviour. Although little studied, it seems likely to draw on hippocampal-centred networks incorporating prefrontal cortical regions such as the middle and inferior frontal gyri (Carter et al., 2006).

Overall, this perspective on the subjective feeling of pain fits with decoding studies showing that pain perception derives from information contained with disparate brain networks (Wager, Atlas, et al., 2013; Woo et al., 2015). The key to understanding pain, therefore, is understanding precisely what information processing operations are being implemented in these specific underlying circuits. Ultimately, this is a necessary step in understanding how brain activity gives rise to pain experience, whether abnormalities in brain computations contribute to chronic pain, and how to identify targets for new pain therapies.

In summary, the empirical evidence suggests there are multiple parallel hierarchical circuits that implement distinct components of information processing - and ultimately yield effective and efficient pain behaviour. This tells us something fundamental about pain: the fact that there are multiple and parallel pathways suggests, from the perspective of a generative or forward model, being in a state of pain has multiple sensory consequences in many domains. These can include nociceptive and interoceptive signals, while at the same time, making predictions in the somatosensory and proprioceptive domain. Furthermore, it tells us that pain - as a construct - must be deep in the hierarchy. In turn, this suggests that pain is quintessentially constructive in nature. On the active inference view, pain can be thought of as a hypothesis entertained by the brain that provides the simplest explanation for the pattern of inputs that the individual is currently experiencing. In virtue of the fact that it is a high-level construct that necessarily entails descending predictions of precision, there is also an argument that it loses phenomenal transparency (Limanowski and Friston, 2018). In other words, 'not only do I have qualitative experience of pain, but I know that I am in pain'.

Credit author statement

BS and FM both conceived, wrote, and edited the manuscript.

Acknowledgements

BS was funded by the Institute of Information Communications Technology Planning Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2019-0-01371, Development of braininspired AI with human-like intelligence), the Wellcome Trust (097490) and Versus Arthritis (21192, 21537). FM was funded by Versus Arthritis (21192) and a Career Development Award from the Medical Research Council (G101584).

References

- Aitchison, L., Lengyel, M., 2017. With or without you: predictive coding and Bayesian inference in the brain. *Curr. Opin. Neurobiol.* 46, 219–227.
- Allen, M., Friston, K.J., 2018. From cognitivism to autopoiesis: towards a computational framework for the embodied mind. *Synthese* 195, 2459–2482.
- Amaral, D.G., Price, J., 1984. Amygdalo-cortical projections in the monkey (*Macaca fascicularis*). *J. Comparative Neurol.* 230, 465–496.
- Anchisi, D., Zanon, M., 2015. A Bayesian perspective on sensory and cognitive integration in pain perception and placebo analgesia. *PLoS ONE* 10 (2), e0117270.
- Apkarian, A.V., Hodge, C.J., 1989. A dorsolateral spinothalamic tract in macaque monkey. *Pain* 37 (3), 323–333.
- Atlas, L.Y., Bolger, N., et al., 2010. Brain mediators of predictive cue effects on perceived pain. *J. Neurosci.* 30 (39), 12964–12977.
- Atlas, L.Y., Doll, B.B., et al., 2016. Instructed knowledge shapes feedback-driven aversive learning in striatum and orbitofrontal cortex, but not the amygdala. *Elife* 5, e15192.
- Basbaum, A.I., Fields, H.L., 1984. Endogenous pain control systems: brainstem spinal pathways and endorphin circuitry. *Annu. Rev. Neurosci.* 7 (1), 309–338.
- Bellman, R., 2013. *Dynamic Programming*. Courier Corporation.
- Berridge, K.C., 2009. Wanting and liking: observations from the neuroscience and psychology laboratory. *Inquiry* 52, 378–398.
- Block, N., 2005. Two neural correlates of consciousness. *Trends Cogn. Sci. (Regul. Ed.)* 9 (2), 46–52.
- Bolles, R.C., 1971. “Species-specific defense reactions. In: *Aversive Conditioning and Learning*. Elsevier, pp. 183–233.
- Bolles, R.C., 1972. The avoidance learning problem. *Psychol. Learn. Motivation* 6, 97–145.
- Borhani, K., Beck, B., Haggard, P., 2017. Choosing, doing, and controlling: implicit sense of agency over somatosensory events. *Psychol. Sci.* 28 (7), 882–893.
- Bornhövd, K., et al., 2002. Painful stimuli evoke different stimulus–response functions in the amygdala, prefrontal, insula and somatosensory cortex: a single-trial fMRI study. *Brain* 125 (6), 1326–1336.
- Botvinick, M., Braver, T., 2015. Motivation and cognitive control: from behavior to neural mechanism. *Annu. Rev. Psychol.* 66.
- Botvinick, M., Toussaint, M., 2012. Planning as inference. *Trends Cogn. Sci. (Regul. Ed.)* 16 (10), 485–488.
- Braid, L., Cahusac, P.M., 2006. Decreased sensitivity to self-inflicted pain. *Pain* 124 (1–2), 134–139.
- Bräscher, A.-K., et al., 2016. Different brain circuitries mediating controllable and uncontrollable pain. *J. Neurosci.* 36 (18), 5013–5025.
- Brown, C.A., Seymour, B., Boyle, Y., et al., 2008a. Modulation of pain ratings by expectation and uncertainty: behavioral characteristics and anticipatory neural correlates. *Pain* 135 (3), 240–250.
- Brown, C.A., Seymour, B., El-Derey, W., et al., 2008b. Confidence in beliefs about pain predicts expectancy effects on pain perception and anticipatory processing in right anterior insula. *Pain* 139 (2), 324–332.
- Bruineberg, J., Kiverstein, J., Rietveld, E., 2018. The anticipating brain is not a scientist: the free-energy principle from an ecological-enactive perspective. *Synthese* 195 (6), 2417–2444.
- Büchel, C., et al., 2014. Placebo analgesia: a predictive coding perspective. *Neuron* 81 (6), 1223–1239.
- Carter, R.M., et al., 2006. Contingency awareness in human aversive conditioning involves the middle frontal gyrus. *Neuroimage* 29 (3), 1007–1012.
- Cauda, F., et al., 2011. Functional connectivity of the insula in the resting brain. *Neuroimage* 55 (1), 8–23.
- Cerminara, N.L., et al., 2009. The periaqueductal grey modulates sensory input to the cerebellum: a role in coping behaviour? *Eur. J. Neurosci.* 29 (11), 2197–2206.
- Colloca, L., Benedetti, F., 2006. How prior experience shapes placebo analgesia. *Pain* 124 (1), 126–133.
- Colombo, M., Seriès, P., 2012. Bayes in the brain—on Bayesian modelling in neuroscience. *The British journal for the philosophy of science* 63 (3), 697–723.
- Craig, A., 2003. A new view of pain as a homeostatic emotion. *Trends Neurosci.* 26 (6), 303–307.
- Craig, A., 2002. How do you feel? Interoception: the sense of the physiological condition of the body. *Nat. Rev. Neurosci.* 3 (8), 655.
- Craig, A., 2009. How do you feel—now? The anterior insula and human awareness. *Nat. Rev. Neurosci.* 10 (1), 59.
- Darian-Smith, C., et al., 1993. Ipsilateral cortical projections to areas 3a, 3b, and 4 in the macaque monkey. *J. Comp. Neurol.* 335 (2), 200–213.
- Daw, N.D., Niv, Y., Dayan, P., 2005. Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat. Neurosci.* 8 (12), 1704.
- Dayan, P., Abbott, L., et al., 2003. *Theoretical neuroscience: computational and mathematical modeling of neural systems*. J. Cogn. Neurosci. 15 (1), 154–155.
- Daw, N.D., O’Doherty, J.P., et al., 2006. Cortical substrates for exploratory decisions in humans. *Nature* 441, 876–879.
- Dayan, P., Abbott, L.F., 2001. *Theoretical Neuroscience*, 806. MIT Press, Cambridge, MA.
- Dayan, P., Daw, N.D., 2008. Decision theory, reinforcement learning, and the brain. *Cognitive, Affective, Behav. Neurosci.* 8 (4), 429–453.
- Dayan, P., Niv, Y., et al., 2006. The misbehavior of value and the discipline of the will. *Neural Netw.* 19 (8), 1153–1160.
- Dayan, P., Sejnowski, T.J., 1996. Exploration bonuses and dual control. *Mach. Learn.* 25 (1), 5–22.
- Delgado, M.R., et al., 2008. The role of the striatum in aversive learning and aversive prediction errors. *Philosophical Trans. R. Soc. London B* 363, 3787–3800 1511.
- Denrell, J., March, J.G., 2001. Adaptation as information restriction: the hot stove effect. *Org. Sci.* 12 (5), 523–538.
- Dickinson, A., Balleine, B., 2002. The role of learning in the operation of motivational systems. *Stevens’ Handbook of Experimental Psychology*.
- Dickinson, A., Dearing, M.F., 1979. Appetitive-aversive interactions and inhibitory processes. *Mech. Learn. Motiv.* 203–231.
- Dimitrova, A., et al., 2003. Cerebellar responses evoked by nociceptive leg withdrawal reflex as revealed by event-related fMRI. *J. Neurophysiol.* 90 (3), 1877–1886.
- Dostrovsky, J., Craig, A., 2013. Ascending projection systems. In: *Wall and Melzack’s Textbook of Pain*, pp. 187–203.
- Ehrsson, H.H., et al., 2007. Threatening a rubber hand that you feel is yours elicits a cortical anxiety response. *Proc. Natl. Acad. Sci.* 104 (23), 9828–9833.
- Ernst, T.M., et al., 2019. The cerebellum is involved in processing of predictions and prediction errors in a fear conditioning paradigm. *BioRxiv*, 600494.
- Erofěeva, M., 1921. *Further Observations Upon Conditioned Reflexes to Nocuous Stimuli*. Bulletin of Institute of Lesgaft.
- Evrard, H.C., Forro, T., Logothetis, N.K., 2012. Von Economo neurons in the anterior insula of the macaque monkey. *Neuron* 74 (3), 482–489.
- Evrard, H.C., Logothetis, N.K., Craig, A., 2014. Modular architectonic organization of the insula in the macaque monkey. *J. Comp. Neurol.* 522 (1), 64–97.
- Evrard, H.C., 2019. The organization of the primate insular cortex. *Front. Neuroanat.* 13, 43.
- Fardo, F., et al., 2017. Expectation violation and attention to pain jointly modulate neural gain in somatosensory cortex. *Neuroimage* 153, 109–121.
- Farkas, K., 2013. Constructing a world for the Senses. *Phenomenal Intention*. 99–115.
- Fields, H.L., 2018. How expectations influence pain. *Pain* 159, S3–S10.
- Flaherty, A.W., Graybiel, A., 1991. Corticostriatal transformations in the primate somatosensory system. Projections from physiologically mapped body-part representations. *J. Neurophysiol.* 66 (4), 1249–1263.
- Flaherty, A.W., Graybiel, A., 1995. Motor and somatosensory corticostriatal projection magnifications in the squirrel monkey. *J. Neurophysiol.* 74 (6), 2638–2648.
- Friston, K., 2005. A theory of cortical responses. *Philosoph. Trans. R. Soc. B* 360, 815–836 1456.
- Friston, K., 2008. Hierarchical models in the brain. *PLoS Comput. Biol.* 4 (11), e1000211.
- Friston, K., 2010. The free-energy principle: a unified brain theory. *Nat. Rev. Neurosci.* 11 (2), 127.
- Friston, K., 2012. The history of the future of the Bayesian brain. *Neuroimage* 62 (2), 1230–1233.
- Friston, K., 2013. Consciousness and hierarchical inference. *Neuropsychanalysis* 15 (1), 38–42.
- Frontera, J.L. et al. (2020). Bidirectional control of fear memories by the cerebellum through the ventrolateral periaqueductal grey. *bioRxiv*.
- Gauriau, C., Bernard, J.-F., 2002. Pain pathways and parabrachial circuits in the rat. *Exp. Physiol.* 87 (2), 251–258.
- Gershman, S.J., 2018. Deconstructing the human algorithms for exploration. *Cognition* 173, 34–42.
- Gershman, S.J., Niv, Y., 2012. Exploring a latent cause theory of classical conditioning. *Learn. Behav.* 40 (3), 255–268.
- Geuter, S., et al., 2017. Functional dissociation of stimulus intensity encoding and predictive coding of pain in the insula. *Elife* 6, e24770.
- Gillan, C.M., Urcelay, G.P., Robbins, T.W., 2016. “An associative account of avoidance. In: *The Wiley Handbook On the Cognitive Neuroscience of Learning*. Wiley Online Library, p. 442.
- Gläscher, J., et al., 2010. States versus rewards: dissociable neural prediction error signals underlying model-based and model-free reinforcement learning. *Neuron* 66 (4), 585–595.
- Grahl, A., Onat, S., Büchel, C., 2018. The periaqueductal gray and Bayesian integration in placebo analgesia. *Elife* 7, e32930.
- Haggard, P., Iannetti, G.D., Longo, M.R., 2013. Spatial sensory organization and body representation in pain perception. *Curr. Biol.* 23 (4), R164–R176.
- Heinricher, M., et al., 2009. Descending control of nociception: specificity, recruitment and plasticity. *Brain Res. Rev.* 60 (1), 214–225.
- Helmchen, C., et al., 2006. Neural activity related to self-versus externally generated painful stimuli reveals distinct differences in the lateral pain system in a parametric fMRI study. *Hum. Brain Mapp.* 27 (9), 755–765.
- Herry, C., Johansen, J.P., 2014. Encoding of fear learning and memory in distributed neuronal circuits. *Nat. Neurosci.* 17 (12), 1644–1654.
- Hird, E.J., et al., 2018. Boundary effects of expectation in human pain perception. *BioRxiv*, 467738. Hoskin, R. et al. (2019). Sensitivity to pain expectations: a Bayesian model of individual differences. *Cognition* 182, 127–139.

- Janak, P.H., Tye, K.M., 2015. From circuits to behaviour in the amygdala. *Nature* 517, 284–293.
- Jepma, M., et al., 2018. Behavioural and neural evidence for self-reinforcing expectancy effects on pain. *Nature Human Behav.* 2 (11), 838–855.
- Johansen, J.P., et al., 2010. Neural substrates for expectation-modulated fear learning in the amygdala and periaqueductal gray. *Nat. Neurosci.* 13 (8), 979–986.
- Kaas, J.H., Stepniowska, L., 2016. Evolution of posterior parietal cortex and parietal-frontal networks for specific actions in primates. *J. Comp. Neurol.* 524 (3), 595–608.
- Kakade, S., Dayan, P., 2002. Dopamine: generalization and bonuses. *Neural Netw.* 15 (4–6), 549–559.
- Kaplan, D.M., 2011. Explanation and description in computational neuroscience. *Synthese* 183 (3), 339.
- Knill, D.C., Pouget, A., 2004. The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends Neurosci.* 27 (12), 712–719.
- Koban, L., et al., 2019. Different brain networks mediate the effects of social and conditioned expectations on pain. *Nat. Commun.* 10 (1), 1–13.
- Kouneiher, F., Charron, S., Koehlin, E., 2009. Motivation and cognitive control in the human prefrontal cortex. *Nat. Neurosci.* 12, 939.
- Lange, I., et al., 2015. The anatomy of fear learning in the cerebellum: a systematic meta-analysis. *Neurosci. Biobehav. Rev.* 59, 83–91.
- Lee, S.W., Shimojo, S., O'Doherty, J.P., 2014. Neural computations underlying arbitration between model-based and model-free learning. *Neuron* 81 (3), 687–699.
- Limanowski, J., Friston, K., 2018. Seeing the dark: grounding phenomenal transparency and opacity in precision estimation for active inference. *Front. Psychol.* 9, 643.
- Litt, M.D., 1988. Self-efficacy and perceived control: cognitive mediators of pain tolerance. *J. Pers. Soc. Psychol.* 54 (1), 149.
- Lobanov, O.V., et al., 2013. Frontoparietal mechanisms supporting attention to location and intensity of painful stimuli. *Pain* 154 (9), 1758–1768.
- Lorenz, J., Minoshima, S., Casey, K., 2003. Keeping pain out of mind: the role of the dorsolateral prefrontal cortex in pain modulation. *Brain* 126 (5), 1079–1091.
- Mackintosh, N.J., 1983. *Conditioning and Associative Learning*. Clarendon Press Oxford.
- Maia, T.V., 2010. Two-factor theory, the actor-critic model, and conditioned avoidance. *Learn. Behav.* 38 (1), 50–67.
- Mancini, F., Bauleo, A., et al., 2014. Whole-body mapping of spatial acuity for pain and touch. *Ann. Neurol.* 75 (6), 917–924.
- Mancini, F., Beaumont, A.-L., et al., 2015. Touch inhibits subcortical and cortical nociceptive responses. *Pain* 156 (10), 1936.
- Mancini, F., Dolgevica, S., et al., 2016. Perceptual learning to discriminate the intensity and spatial location of nociceptive stimuli. *Sci. Rep.* 6.
- Mancini, F., Haggard, P., et al., 2012. Fine-grained nociceptive maps in primary somatosensory cortex. *J. Neurosci.* 32 (48), 17155–17162.
- Marr, D., Poggio, T., 1979. A computational theory of human stereo vision. *Proc. R. Soc. London B* 204, 301–328.
- Melzack, R., 1968. Sensory, motivational, and central control determinants of pain. *The skin Senses*.
- Mohr, C., et al., 2005. The anterior cingulate cortex contains distinct areas dissociating external from self-administered painful stimulation: a parametric fMRI study. *Pain* 114 (3), 347–357.
- Moseley, G.L., Arntz, A., 2007. The context of a noxious stimulus affects the pain it evokes. *Pain* 133 (1–3), 64–71.
- Moulton, E.A., et al., 2010. The cerebellum and pain: passive integrator or active participant. *Brain Res Rev* 65 (1), 14–27.
- Moutoussis, M., Bentall, R.P., et al., 2008. A temporal difference account of avoidance learning. *Network* 19 (2), 137–160.
- Moutoussis, M., Fearon, P., et al., 2014. Bayesian inferences about the self (and others): a review. *Conscious Cogn* 25, 67–76.
- Mower, O. (1960). *Learning theory and behavior*.
- O'Doherty, J.P., Lee, S.W., McNamee, D., 2015. The structure of reinforcement-learning mechanisms in the human brain. *Curr. Opin. Behav. Sci.* 1, 94–100.
- Ongaro, G., Kaptchuk, T.J., 2019. Symptom perception, placebo effects, and the Bayesian brain. *Pain* 160 (1), 1–4.
- Oshiro, Y., et al., 2009. Brain mechanisms supporting discrimination of sensory features of pain: a new model. *J. Neurosci.* 29 (47), 14924–14931.
- Owens, A.P., et al., 2018. Interoceptive inference: from computational neuroscience to clinic. *Neurosci. Biobehav. Rev.* 90, 174–183.
- Pandya, D.N., Vignolo, L.A., 1969. Interhemispheric projections of the parietal lobe in the rhesus monkey. *Brain Res.* 15 (1), 49–65.
- Parr, T., Friston, K.J., 2018. Generalised free energy and active inference: can the future cause the past? *BioRxiv* 304782.
- Pezzulo, G., Rigoli, F., Friston, K., 2015. Active Inference, homeostatic regulation and adaptive behavioural control. *Prog. Neurobiol.* 134, 17–35.
- Ploghaus, A., et al., 1999. Dissociating pain from its anticipation in the human brain. *Science* 284, 1979–1981.
- Pons, T., Kaas, J., 1986. Corticocortical connections of area 2 of somatosensory cortex in macaque monkeys: a correlative anatomical and electrophysiological study. *J. Comp. Neurol.* 248 (3), 313–335.
- Price, D.D., Greenspan, J.D., Dubner, R., 2003. Neurons involved in the exteroceptive function of pain. *Pain* 106 (3), 215–219.
- Ramstead, M.J., et al., 2019. Multiscale integration: beyond internalism and externalism. *Synthese* 1–30.
- Rao, R.P., Ballard, D.H., 1999. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* 2 (1), 79.
- Robbe, D., 2018. To move or to sense? Incorporating somatosensory representation into striatal functions. *Curr. Opin. Neurobiol.* 52, 123–130.
- Roeder, Z., et al., 2016. The parabrachial complex links pain transmission to descending pain modulation. *Pain* 157 (12), 2697.
- Roy, M., et al., 2014. Representation of aversive prediction errors in the human periaqueductal gray. *Nat. Neurosci.* 17 (11), 1607–1612.
- Salomons, T.V., et al., 2007. Individual differences in the effects of perceived controllability on pain perception: critical role of the prefrontal cortex. *J. Cogn. Neurosci.* 19 (6), 993–1003.
- Schultz, W., 2016. Dopamine reward prediction-error signalling: a two-component response. *Nat. Rev. Neurosci.* 17 (3), 183.
- Schultz, W., Dayan, P., Montague, P.R., 1997. A neural substrate of prediction and reward. *Science* 275, 1593–1599.
- Seminowicz, D.A., Moayed, M., 2017. The dorsolateral prefrontal cortex in acute and chronic pain. *J. Pain* 18 (9), 1027–1035.
- Seth, A.K., 2013. Interoceptive inference, emotion, and the embodied self. *Trends Cogn. Sci. (Regul. Ed.)* 17 (11), 565–573.
- Seymour, B., Dolan, R., 2013. "Emotion, motivation, and pain. In: Wall & Melzack's Textbook of Pain. Elsevier Health Sciences, Philadelphia, pp. 248–255.
- Seymour, B., 2019. Pain: a precision signal for reinforcement learning and control. *Neuron* 101 (6), 1029–1041.
- Seymour, B., Daw, N.D., et al., 2012. Serotonin selectively modulates reward value in human decision-making. *J. Neurosci.* 32 (17), 5833–5842.
- Seymour, B., O'Doherty, J.P., Koltzenburg, M., et al., 2005. Opponent appetitive-aversive neural processes underlie predictive learning of pain relief. *Nat. Neurosci.* 8 (9), 1234–1240.
- Siegel, P., Wepsic, J.G., 1974. Alteration of nociception by stimulation of cerebellar structures in the monkey. *Physiol. Behav.* 13 (2), 189–194.
- Snyder, R.L., Faull, R.L., Mehler, W.R., 1978. A comparative study of the neurons of origin of the spinocerebellar afferents in the rat, cat and squirrel monkey based on the retrograde transport of horseradish peroxidase. *J. Comp. Neurol.* 181 (4), 833–852.
- Stevens, R.T., London, S.M., Apkarian, A.V., 1993. Spinothalamic cortical projections to the secondary somatosensory cortex (SII) in squirrel monkey. *Brain Res.* 631 (2), 241–246.
- Sutton, R.S., Barto, A.G., 1981. Toward a modern theory of adaptive networks: expectation and prediction. *Psychol. Rev.* 88 (2), 135.
- Sutton, R.S., Barto, A.G., 1998. *Reinforcement learning: An introduction*, Vol. 1. MIT press, Cambridge.
- Tabor, A., Burr, C., 2019. Bayesian learning models of pain: a call to action. *Current Opinion in Behavioral Sciences* 26, 54–61.
- Tabor, A., et al., 2017. Pain: a statistical account. *PLoS Comput. Biol.* 13 (1), e1005142.
- Taylor, V.A., et al., 2017. Learned expectations and uncertainty facilitate pain during classical conditioning. *Pain* 158 (8), 1528–1537.
- Seymour, B., O'Doherty, J.P., Dayan, P., et al., 2004. Temporal difference models describe higher-order learning in humans. *Nature* 429 (6992), 664–667.
- Thrun, S.B. (1992). *Efficient exploration in reinforcement learning*.
- Todd, A.J., 2010. Neuronal circuitry for pain processing in the dorsal horn. *Nat. Rev. Neurosci.* 11 (12), 823.
- Tolman, E.C., 1948. Cognitive maps in rats and men. *Psychol. Rev.* 55 (4), 189.
- Wager, T.D. (2005). *Expectations and anxiety as mediators of placebo effects in pain*.
- Vierck, C.J., et al., 2013. Role of primary somatosensory cortex in the coding of pain. *Pain* 154 (3), 334–344.
- Wager, T.D., Atlas, L.Y., et al., 2013. An fMRI-based neurologic signature of physical pain. *N. Engl. J. Med.* 368 (15), 1388–1397.
- Wager, T.D., Rilling, J.K., et al., 2004. Placebo-induced changes in fMRI in the anticipation and experience of pain. *Science* 303, 1162–1167.
- Wald, A., 1947. An essentially complete class of admissible decision functions. *Ann. Math. Statist.* 18 (4), 549–555.
- Wang, O., et al., 2018. Model-based and model-free pain avoidance learning. *Brain Neurosci. Adv.* 2, 2398212818772964.
- Watson, A., et al., 2006. Categories of placebo response in the absence of site-specific expectation of analgesia. *Pain* 126, 115–122.
- Wiech, K., 2016. Deconstructing the sensation of pain: the influence of cognitive processes on pain perception. *Science* 354, 584–587.
- Wiech, K., Kalisch, R., et al., 2006. Anterolateral prefrontal cortex mediates the analgesic effect of expected and perceived control over pain. *J. Neurosci.* 26 (44), 11501–11509.
- Wiech, K., Vandekerckhove, J., et al., 2014. Influence of prior information on pain involves biased perceptual decision-making. *Curr. Biol.* 24 (15), R679–R681.
- Wilson, R.C., et al., 2014. Humans use directed and random exploration to solve the explore–exploit dilemma. *J. Experimen. Psychol.* 143 (6), 2074.
- Wittmann, B.C., et al., 2008. Striatal activity underlies novelty-based choice in humans. *Neuron* 58 (6), 967–973.
- Woo, C.-W., et al., 2015. Distinct brain systems mediate the effects of nociceptive input and self-regulation on pain. *PLoS Biol.* 13 (1), e1002036.
- Yoshida, W., et al., 2013. Uncertainty increases pain: evidence for a novel mechanism of pain modulation involving the periaqueductal gray. *J. Neurosci.* 33 (13), 5638–5646.
- Yu, A.J., Dayan, P., 2005. Uncertainty, neuromodulation, and attention. *Neuron* 46 (4), 681–692.
- Zhang, S., Mano, H., Ganesh, G., et al., 2016. Dissociable learning processes underlie human pain conditioning. *Curr. Biol.* 26 (1), 52–58.
- Zhang, S., Mano, H., Lee, M., et al., 2018. The control of tonic pain by active relief learning. *Elife* 7, e31949.