

Research



Cite this article: Parr T, Friston KJ. 2017
Uncertainty, epistemics and active inference.
J. R. Soc. Interface **14**: 20170376.
<http://dx.doi.org/10.1098/rsif.2017.0376>

Received: 22 May 2017

Accepted: 27 October 2017

Subject Category:

Life Sciences—Mathematics interface

Subject Areas:

computational biology, systems biology

Keywords:

uncertainty, Bayesian, active inference,
neuromodulation, acetylcholine, noradrenergic

Author for correspondence:

Thomas Parr

e-mail: thomas.parr.12@ucl.ac.uk

Uncertainty, epistemics and active inference

Thomas Parr and Karl J. Friston

Wellcome Trust Centre for Neuroimaging, Institute of Neurology, University College London, 12 Queen Square, London WC1N 3BG, UK

TP, 0000-0001-5108-5743

Biological systems—like ourselves—are constantly faced with uncertainty. Despite noisy sensory data, and volatile environments, creatures appear to actively maintain their integrity. To account for this remarkable ability to make optimal decisions in the face of a capricious world, we propose a generative model that represents the beliefs an agent might possess about their own uncertainty. By simulating a noisy and volatile environment, we demonstrate how uncertainty influences optimal epistemic (visual) foraging. In our simulations, saccades were deployed less frequently to regions with a lower sensory precision, while a greater volatility led to a shorter inhibition of return. These simulations illustrate a principled explanation for some cardinal aspects of visual foraging—and allow us to propose a correspondence between the representation of uncertainty and ascending neuromodulatory systems, complementing that suggested by Yu & Dayan (Yu & Dayan 2005 *Neuron* **46**, 681–692. (doi:10.1016/j.neuron.2005.04.026)).

1. Introduction

In this paper, we address the computational basis for the representation of uncertainty by the brain, and its consequences for epistemic (information gathering¹) behaviour. We focus on two sources of uncertainty: uncertainty concerning the temporal evolution of environmental states, and uncertainty about the mapping from (hidden) states of the world to sensory observations. Both may arise either from uncertainty inherent in the external world or from noise in neuronal signalling [1]. The first source of uncertainty corresponds to the volatility of state transitions, while the second corresponds to sensory noise and ambiguity. The latter has previously been addressed in the context of predictive coding, in which sensory precision (i.e. inverse variance) modulates the (possibly attentional) gain of ascending prediction errors [2]. This modulatory effect is a direct consequence of (Bayes) optimal evidence accumulation (cf. the Kalman gain of Bayesian filters in engineering). This formulation of attention appeals to the notion of the brain as a statistical organ: an organ that infers the causes of its sensations using internal models of how sensory impressions are generated by continuous states of the world [3,4]. Here, we consider the role of precision in discrete state space models.

Both predictive coding and the (Bayesian) decision processes described in this paper conform to the free energy principle. This principle states that, to prevent its entropy (time average of surprise) growing indefinitely, an agent must maintain an upper bound on surprise [5]. This upper bound is the variational *free energy* of sensory samples [6,7]. The free energy is a function of beliefs about hidden (latent) states, s , and a function of observations, o_{τ} , defined as

$$F = -E_{Q(s)}[\ln P(\tilde{o}, s) - \ln Q(s)].$$

The notation \tilde{o} indicates a trajectory of observations through time; i.e. $[o_1, o_2, \dots, o_T]^T$. As written, this equation is very general. To establish its consequences in concrete scenarios, it is necessary to specify a particular form of the

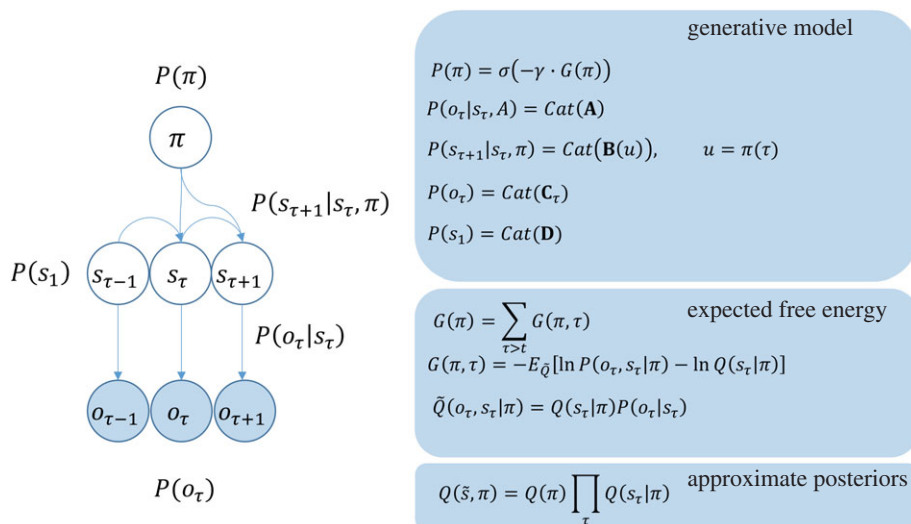


Figure 1. MDP: this Bayesian network *on the left* illustrates the conditional dependencies, and independencies, between the variables in the generative model (see the main text for a description of the variables). The panels *on the right* give the forms of the distributions in the generative model, in addition to defining the expected free energy, and specifying the (mean-field) factorization of the approximate posterior distributions (beliefs) the agent possesses. (Online version in colour.)

generative model, $P(\tilde{o}, s)$, and a factorization of the beliefs, $Q(s)$, an agent possesses. After the form of the generative model and beliefs have been specified, standard schemes can be used to simulate perception by changing beliefs about hidden or latent states of the world to minimize free energy. This usually involves variational filtering (i.e. Bayesian filtering) for continuous state space models and variational message passing (i.e. belief propagation) for discrete state space models. To simulate action and perception, both beliefs and action minimize free energy. This is known as *active inference*. The novel theoretical contribution of this paper is the inclusion of explicit beliefs about uncertainty in a discrete state space generative model.

We show through simulation that epistemic foraging is heavily influenced by the beliefs an agent has about the volatility and sensory precision of their environment [8]. The temporal dynamics of visual search, including the phenomenon of ‘inhibition of return’, follow naturally from this formulation. The formal contribution of sensory precision to salience dissolves the ‘dark room problem’ [9] associated with active inference, without needing to invoke additional prior beliefs [10]. In the following, we will briefly overview the structure of the generative model we have used previously [11–17]. This model is then supplemented with volatility and precision parameters, to illustrate their role through simulations of visual foraging. These simulations are used to build an intuition behind the phenomenology of Bayes-optimal searches under uncertainty. In the final section, we turn to neuronal implementations of the ensuing ‘precision engineered’ message passing and consider the implications for—and predictions of—empirical studies of attention and neurotransmitter function.

2. Markov decision process

A Markov decision process (MDP) is a form of probabilistic generative model [18], defined in a discrete state space. The latent variables of an MDP are hidden states s_τ and policies, π . The conditional dependencies in the model are expressed graphically in figure 1. Hidden states, s_τ , generate observable sensory

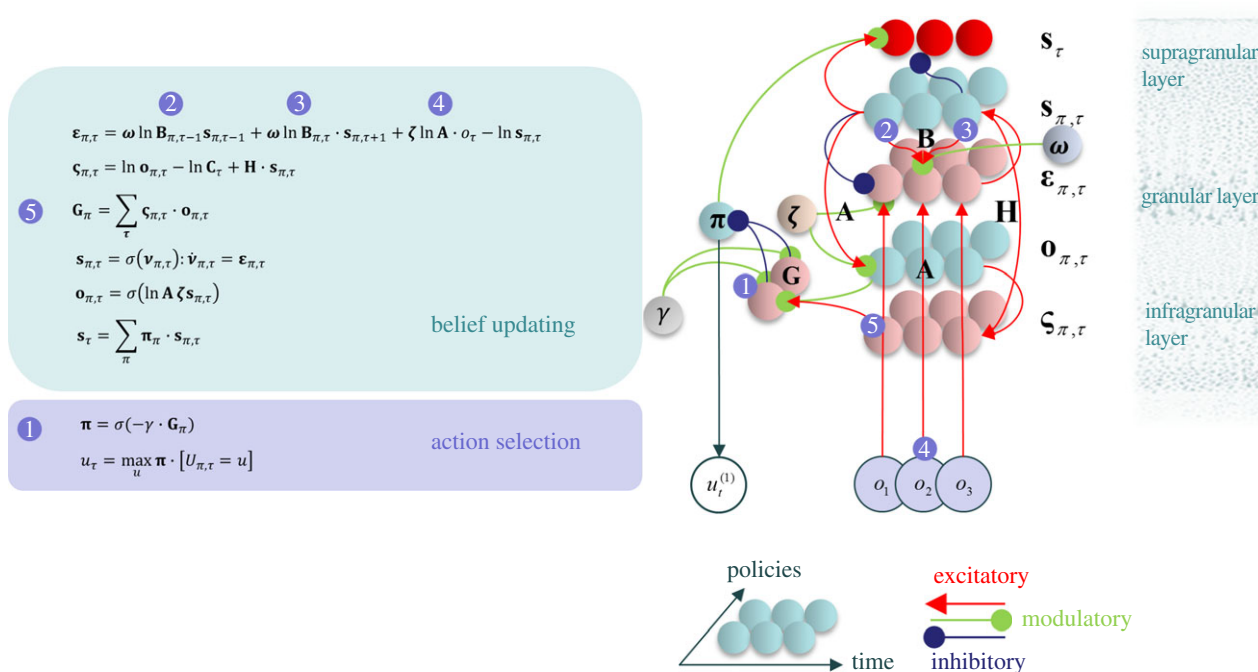
data, o_τ , with probabilities expressed in a likelihood matrix $P(o_\tau = i | s_\tau = j) = \mathbf{A}_{ij}$. It is this matrix that allows for ‘top-down’ predictions analogous to those in many descriptions of perception [3,4,19,20]. The states evolve through time according to a transition probability matrix, $P(s_{\tau+1} = i | s_\tau = j, \pi) = \mathbf{B}(u)_{ij}$, so they depend only on the state at the previous time, and on the policy, π , pursued by an agent. It is important to note that, although the first dependency renders the process ‘Markovian’, the dependence on policies of arbitrary lengths breaks this property. In other words, this form of generative model can account for processes with ‘memory’, something that is further enhanced in hierarchical extensions of this model [16,21]. Policies represent sequences of actions, $u = \pi(\tau)$, that determine the form of the state transition matrix. To complete the specification of the model for active inference, it is necessary to introduce prior distributions over outcomes, $P(o_\tau = i) = \mathbf{C}_\tau$, the initial hidden state, $P(s_1 = i) = \mathbf{D}_i$, and the policy. Active inference treats policy selection as a Bayesian model selection problem. In other words, policies are selected based upon the free energy, $G(\pi)$ expected on pursuing that policy in question [13]. This quantity is defined in figure 1. Using a softmax (normalized exponential) function to convert the expected free energy to a probability distribution, we can write a prior belief over policies as

$$P(\pi) = \sigma(-\gamma \cdot G(\pi)).$$

In this equation, γ is an inverse temperature parameter, that corresponds to the confidence (or precision) of beliefs about policies. This has been used extensively in previous formulations of active inference, and has been proposed as a computational homologue of dopamine signalling [12,14,15].

3. Precision and volatility

To equip the agent with beliefs about the uncertainty in both the transitions of hidden states (i.e. state precision) and the likelihood mapping from hidden states to outcomes (i.e. sensory precision), we now introduce precision parameters to the agent’s generative model. These are inverse temperature parameters, analogous to γ used for the policy prior. We first



augment the likelihood distribution with a sensory precision, ζ :

$$P(o_\tau = i | s_\tau = j, \zeta^{(j)}) = \frac{\mathbf{A}_{ij}^{\zeta^{(j)}}}{\sum_k \mathbf{A}_{kj}^{\zeta^{(j)}}}.$$

This is a Gibbs measure, commonly expressed as a softmax function [22], for which the denominator is a normalizing constant (partition function). In this equation, ζ is the analogue of precision in predictive coding formulations of attentional gain [2]. Note that each value s_τ can take is associated with its own precision. When $\zeta^{(j)} = 1, \forall j$, this equation reduces to the formulation in figure 1. The same approach can be followed to define the precision of state transitions (i.e. inverse volatility), ω :

$$P(s_{\tau+1} = i | s_\tau = j, \pi, \omega) = \frac{\mathbf{B}(u)_{ij}^\omega}{\sum_k \mathbf{B}(u)_{kj}^\omega}.$$

It is simple to extend this, so that $\omega^{(j)}$ is different for each s_τ as $\zeta^{(j)}$ is, but this is not necessary for the simulations that follow. For simplicity, we will assume $\omega^{(j)} = \omega, \forall j$ for the remainder of this paper.

Given the structure of the MDP shown in figure 1, and the precision parameters defined here, it is possible to express the free energy explicitly, and to find its minimum with respect to each factor in the approximate posterior distribution. In doing so (appendix A), we arrive at the belief update equations in figure 2 for expected states of the world, and expectations about the policies currently being pursued. If these belief update or belief propagation equations are interpreted in terms of message passing in the brain, the resulting connectivity closely resembles that of a cortical column, that participates in a cortico-subcortical loop. Figure 2 illustrates this correspondence by expressing the belief updates for expected states, outcomes and policies ($s_{\pi,\tau}, \mathbf{o}_{\pi,\tau}, \pi$), in terms of auxiliary variables; namely, prediction errors and log expectations ($\varepsilon_{\pi,\tau}, \mathbf{v}_{\pi,\tau}$) that play the role of neuronal depolarization.

4. Simulations

To illustrate the influence of beliefs about uncertainty on behaviour, the generative model of figure 3 was used to simulate epistemic foraging. The generative model includes four stimuli, whose identity can change stochastically. These stimuli are mapped, noisily, to observable outcomes. Each stimulus is associated with a hidden state that defines its identity. An additional hidden state is the current eye position that determines which of the stimuli is observed. This is associated with an identity mapping to a proprioceptive outcome indicating the current eye position, in a manner consistent with previous MDP models of saccadic eye movements [17]. In brief, this means that given the hidden states (namely, where the agent looks and the states of the stimulus at that position), one can generate probabilistic outcomes (proprioceptive information about where the agent is looking and exteroceptive outcomes reporting stimulus identity).

The behaviour observed in the simulations can be explained by referring to the agent's beliefs about policies. As shown in figure 1, a system that engages in active inference selects policies with a low expected free energy. A rearrangement of the equation for expected free energy at a future time gives the following

$$G(\pi, \tau) = -E_{\tilde{Q}}[\ln P(o_\tau)] - E_{Q(o_\tau, \pi)}[D_{KL}[Q(s_\tau | o_\tau, \pi) || Q(s_\tau | \pi)]].$$

The first term here corresponds to prior preferences. These are uniform across all outcomes in these simulations. The second term corresponds to epistemic value or salience, and it is this that drives the active sampling of a visual scene [23,24]. The greater the change expected in beliefs with and without future outcomes, the lower the expected free energy. In other words, the salience reflects the expected information gain or resolution of uncertainty about states of the world. For a location associated with a low sensory precision, an observation is unlikely to elicit a substantial change in the

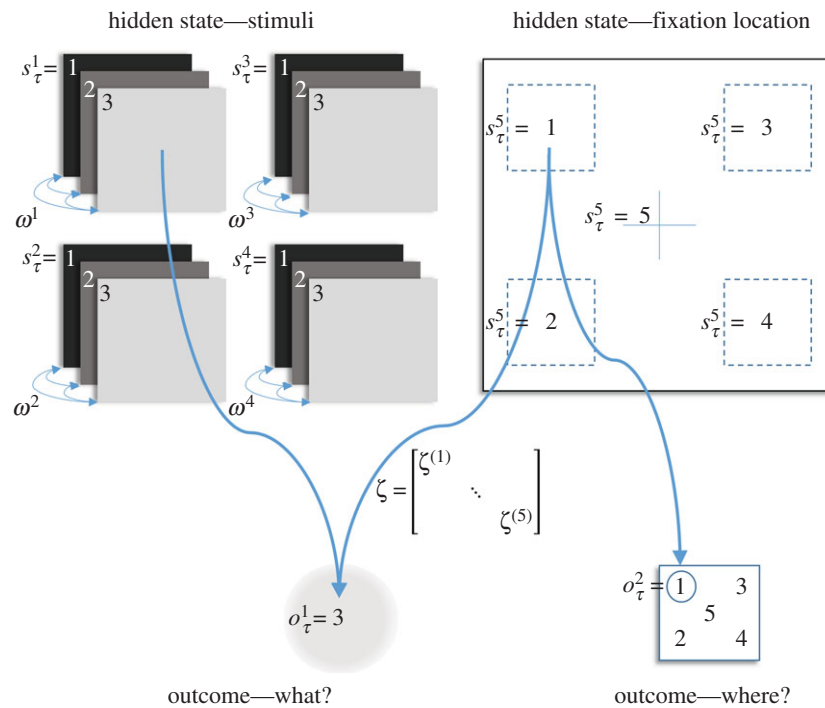


Figure 3. Generative model: the generative model used for the simulations involves five sorts of hidden states. These are four stimuli, and the current eye position. Each stimulus can take one of three identities (illustrated above as dark, medium or light shades). There are five possible fixation locations that correspond to the locations of each stimulus, and the fixation cross. These hidden states map to two outcome modalities: *what* (stimulus identity) and *where* (location). The mapping to the location (i.e. proprioceptive) outcome is an identity mapping from the hidden state representing fixation locations. Each fixation location is associated with a different sensory precision (diagonal elements of ζ), mapping to the *what* (i.e. exteroceptive) outcome. Precisions associated with all other hidden states are 1. Each stimulus hidden state can change identity with a transition precision ω . A volatility of $\omega^{-1} = 0$ would correspond to an identity transition matrix. Crucially, the prior beliefs about outcomes are uniform, so there are no prior preferences in play. All behaviour resulting from this generative model is therefore purely epistemic and uncertainty reducing in nature. (Online version in colour.)

posterior, so a saccade to such a location is less likely to be selected. This behaviour is illustrated in figure 4a, which shows a simulated sequence of actions (i.e. saccadic eye movements) over eight sensory samples.

Heuristically, this is why a well-lit room, with precise sensory information, is preferable to a dark room; i.e. precise sensory cues that resolve ambiguity have greater epistemic affordance and are more likely to be sampled. More colloquially, this sort of behaviour recapitulates the joke about the drunkard looking for a lost key under a streetlamp (the ‘Street-light effect’) [25]. Notably, the drunkards ‘cognitive bias’ is entirely Bayes optimal on an active inference view. The greater frequency of saccades to stimuli with a higher volatility (figure 4b) can be similarly explained. On making an observation concerning a stimulus, the expected posterior at the next time step should be very close to the current belief, for that stimulus. Recent observations are thus associated with a lower salience that then gradually increases over time, as the probability that the hidden state has transitioned to a new value increases (figure 5). In other words, knowing the state of a stimulus in a particular location means there is no further information to be gained by sampling that location and it loses its salience. Note that salience is an attribute of both the world and the agent’s beliefs about the world. However, if the stimulus can change, the salience of its location will increase slowly over time with uncertainty about its current status.

This phenomenon is consistent with the ‘forgetting slopes’ determined by calculating the error in reports about a stimulus at different times following presentation [29], and with theoretical analyses of the properties of the attractor networks

used to explain the maintenance of working memory signals [30]. The concept of ‘inhibition of return’ [31,32] naturally emerges from this formulation, as an agent becomes less likely to return to the same location for a temporally limited period following a fixation. A stimulus with high volatility (figure 4b) would then have a shorter inhibition of return, leading to a greater frequency of fixation. Formulating volatility in this way means that the ω parameter can be estimated from real subjects simply by measuring the length of the inhibition of return.

In summary, using a very simple but plausible formulation of active inference in the context of searching a simple visual scene, we find a natural explanation for two key phenomena in visual search; namely the attractiveness of salient, uncertainty reducing target locations and inhibition of return that depends upon the volatility of a visual scene. Crucially, both of these phenomena rest on encoding the uncertainty or precision of state transitions and the generation of (visual) outcomes from hidden states. In what follows, we now consider the neuronal encoding of precision and the intimate relationship between salience, attention and epistemic affordance.

5. The neurobiology of precision

We have presented some elementary simulations to illustrate fairly straightforward phenomena that emerge under active (Bayesian) inference; namely, that the encoding of uncertainty or precision nuances perception and action in a




Figure 4. Simulated visual foraging. These plots show the simulated eye tracking data (19 saccades) for the MDP, and the influence of varying sensory precision and volatility. (a) The situation in which precision is (believed to be) low when the top left location is foveated is shown. (b) The precisions for each location are equal, as are the volatilities for each stimulus. (c) The lower right stimulus is associated with increased volatility. (Online version in colour.)

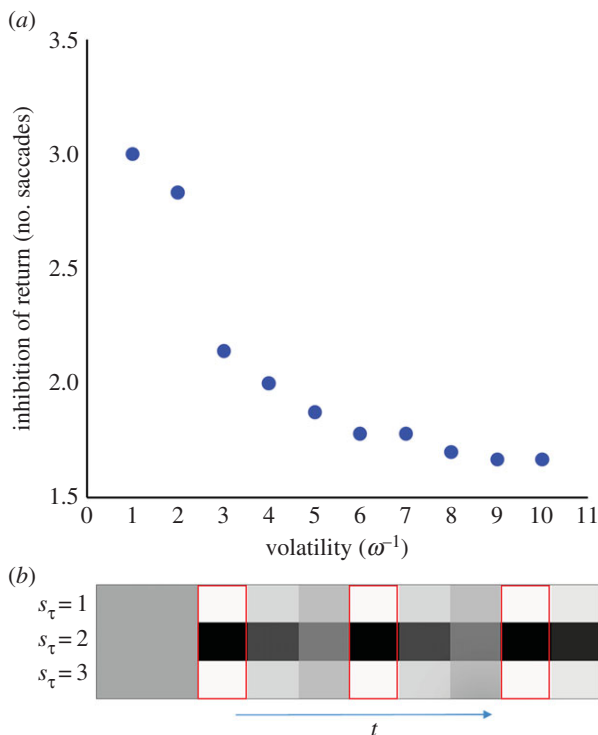


Figure 5. Simulated neuronal encoding and inhibition of return. (a) The influence of volatility beliefs on the simulated inhibition of return (quantified by the average number of saccades between fixations) for one of the stimulus locations. This is consistent with empirical data using noradrenergic pharmacological manipulations [26], and pupillometry [27,28]. (b) The simulated neuronal encoding over time for three units, each representing a different hypothesis about the identity (1, 2 or 3) of one of the hidden states. Darker shades indicate greater firing rates. The red boxes indicate the times during which the location corresponding to this hidden state is foveated. (Online version in colour.)

fundamental but sensible way. In what follows, we revisit the encoding of uncertainty from the point of view of plausible neuromodulatory transmitter systems, the implicit computational anatomy and implications for neuropsychology. This treatment is not a review of the psychology of prediction under uncertainty but a selective survey of empirical findings that speak to the computational architecture of the preceding sections.

5.1. The pharmacology of uncertainty

While we have contrasted sensory uncertainty with volatility, other authors have emphasized the difference between ‘expected uncertainty’ and ‘unexpected uncertainty’ [33]. These two dimensions of classification align more closely than might initially appear. ‘Expected uncertainty’ is framed in terms of a belief that a cue has low validity. In other words, the mapping from a task context (hidden states) to an informative observation is believed to be imprecise. This exactly corresponds to imprecision ζ^{-1} . ‘Unexpected uncertainty’ is described in terms of stochastic changes to the task context. Such changes are necessarily determined by a probability transition matrix, so this form of uncertainty is associated with ω^{-1} . Just as γ has been associated with dopaminergic signalling [15], these two forms of uncertainty correspond to the activity in neuromodulatory systems; specifically, the cholinergic [34,35] and noradrenergic systems [36].

Nicotinic acetylcholine receptors are expressed presynaptically at thalamocortical synapses in cortical layers 3 and 4 [37,38], the laminar targets of first-order thalamic nuclei [39]. These have been shown to modulate sensory gain [40,41] in the visual system. This is consistent with ζ in both the anatomy shown in figure 2, in which it modulates the synapses carrying sensory data to the cortex, and in the functional role implied by the update equations. Acetylcholine receptors are also found in other cortical layers [42,43], including deeper layers, which is again consistent with figure 2. Muscarinic receptors also appear to have an important role in attentional gain mechanisms [44].

Noradrenaline [33,36] has been proposed to signal uncertainty about state transitions. It has also been implicated in modulation of the balance between exploitative and exploratory (epistemic) behaviours [45]. Our simulations demonstrate the consistency of these propositions, as a greater transition uncertainty means that the salience or epistemic affordance of a particular sampling of the environment increases at a faster rate in a volatile context. Pupillary data, associated with catecholamine signalling [46], provide additional support to the hypothesis that noradrenaline is involved in signalling ω , as dilatation occurs during the delay period of working memory tasks [47]. Such tasks require the maintenance of beliefs about a given stimulus identity, that correspond to the belief that the stimulus identity will not change throughout the delay

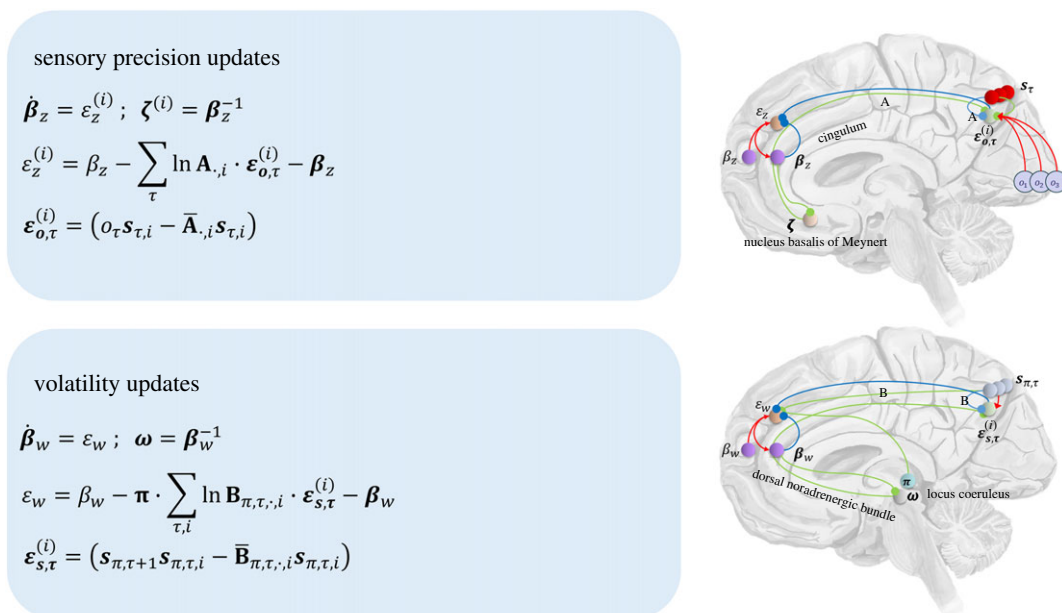


Figure 6. Uncertainty updates and neuroanatomy. *On the left*, the Bayesian update equations for the representations of sensory precision and volatility are shown. These are derived in appendix A, and are presented here as gradient ascents. *On the right*, the variables in these equations are assigned to neuronal populations. This assignment is speculative, but is intended to show how the updates could be implemented in the brain. It is also consistent with the observations in the main text that ascending neuromodulatory systems are likely to signal changes in precision or volatility to the cortex, and that prefrontal regions are well placed to compute the necessary signals. (Online version in colour.)

period (i.e. that the precision of state transitions ω is high). The relationship between pupillary diameter and central noradrenaline offers a means to test the predictions of this neurochemical model of precision signalling. By manipulating the volatility of stimulus transitions, we can test the hypothesis that increases in ω are associated with pupillary constrictions, and that increases in ω are accompanied by dilatations. Crucially, we would not expect modulations in other precision parameters to induce these changes.

5.2. The neuroanatomy of uncertainty

Having parametrized both sensory precision and volatility, we are now in a position to derive Bayes optimal updates for these parameters (see appendix A). This means that the agent can infer the precision of its environment, in terms of both likelihood mappings and state transitions. The resulting update equations are shown in figure 6. This figure additionally shows how the Bayesian updates for ζ and ω could map onto the connectivity between the cortex and the noradrenergic and cholinergic systems. These are related to cortical areas via the cingulum, and the dorsal noradrenergic bundle. Damage to the latter has been linked to deficits in epistemic behaviour [48,49] and attentional set-shifting [50]. Disruption of the dorsal noradrenergic bundle has also been associated with impaired extinction of a conditioned stimulus [51], perhaps reflecting a representation of very low volatility.

If volatility is signalled by noradrenaline, the networks computing this quantity should interact with the locus coeruleus, a noradrenergic brainstem nucleus that projects too much of the cortex [52]. Anterograde tracing has demonstrated that the prefrontal cortex is a source of projections to the locus coeruleus [53]. Pharmacological manipulations [54] show that these projections influence the activity of brainstem

noradrenergic neurons. Specifically, inactivation of frontal regions causes a sustained increase in locus coeruleus firing. This makes these regions good candidate sites for the computation of volatility. In figure 6, volatility is determined using volatility prediction errors in the prefrontal cortex, making use of projections from sensory areas (that carry state prediction errors). The locus coeruleus then sends a precision (inverse volatility) signal to the cortex. Given the close association between central noradrenaline and pupillary diameter [46], the dynamics of the Bayesian updates given here could be incorporated into an MDP-based generative model of pupillary data, first to establish the validity of the updates as a description of noradrenergic signalling, and then as part of an generative model of empirical responses that can be elicited experimentally [55].

Prefrontal regions also project to the basal forebrain [56], the primary source of cholinergic projections to the cortex. To reach the cortex, fibres from the basal forebrain pass the corpus callosum rostrally, before joining the cingulum [43]. Cholinergic axons leave this white matter bundle to diffusely innervate the cerebrum. Figure 6 shows that sensory precision can be calculated in a manner analogous to volatility. Outcome prediction errors in sensory areas are propagated to frontal regions that calculate a precision-related prediction error. This is used to generate a signal to the nucleus basalis of Meynert; a forebrain nucleus that provides a cholinergic signal to sensory cortices.

5.3. The neuropsychology of uncertainty

The disruption of the dopaminergic modulation of policy precision can result in disease states, including Parkinson's disease [15,57,58], and schizophrenia [59–61]. Similarly, the neurotransmitter systems associated here with sensory precision and volatility are disrupted in a range of neuropsychiatric

Table 1. Summary of the functional anatomy of precision.

precision	functional role	neurotransmitter system	neuroanatomy
ζ	encoding the precision of outcomes given hidden states (cf. attention and expected uncertainty [34,35])	cholinergic	nucleus basalis of Meynert
ω	encoding the precision of state transitions (cf. volatility and unexpected uncertainty [36])	noradrenergic	locus coeruleus
γ	Encoding the precision of beliefs about policies (cf. action selection [15])	dopaminergic	substantia nigra pars compacta, ventral tegmental area

disorders. Depletion of acetylcholine is associated with Alzheimer's disease [62,63], while disruptions of noradrenaline signalling are thought to contribute to anxiety [64], post-traumatic stress disorder [65], depression [66] and Wernicke–Korsakoff encephalopathy [67,68]. Additionally, the lateral asymmetry of noradrenergic projections in the fore-brain [69], reflected in pupillary responses [70], hints at a role in hemineglect [71]. A formal description of the computational processes that are disrupted in these disorders allows for the development of a computational phenotyping [55] of patients. This may aid in the characterization of defective neurophysiology, making use of the process theory [12] associated with active inference.

6. Discussion

In order to act optimally in an ambiguous and volatile world, it is necessary to possess a generative model that incorporates key forms of uncertainty. Through their influence on salience—a component of the expected free energy of a policy—uncertainty can have a profound influence on epistemic behaviour. For example, volatile contingencies call for a shorter inhibition of return, while imprecise sensory information is inherently less epistemically valuable. The parametrization in this paper facilitates the derivation of Bayesian updates for precision parameters—that could be implemented by a network of sensory, prefrontal and subcortical structures. Functionally and anatomically, the ascending cholinergic and noradrenergic systems are plausible neurobiological substrates for the computational processes described in this paper.

The story on offer here provides a coherent and formal account of neuromodulation in the brain that is broadly consistent with previous neurobiological accounts of perception and decision-making [8,72,73]. In brief, there are three fundamental sorts of beliefs that determine behaviour: (i) beliefs about outcomes given hidden or latent states of the world, (ii) beliefs about states of the world, and (iii) beliefs about policies given states of the world. Each of these sets of beliefs is equipped with an uncertainty or precision that may be encoded by specific modulatory neurotransmitter systems. The evidence reviewed above—and in [8,33]—speaks to the following (summarized in table 1): (i) cholinergic systems encode the precision of beliefs about outcomes given states of the world (cf. attention and expected uncertainty); (ii) noradrenergic systems encode the precision of state transitions (cf. volatility and unexpected uncertainty); and (iii)

dopaminergic systems encode the precision of beliefs about policies (cf. action selection). The coherent aspect of this account rests on the fact that all three systems play the same computational role; namely, an encoding of precision. Furthermore, all three neurotransmitter systems have the same basic effects on synaptic transmission; namely, a neuromodulatory gain control.

Data accessibility. The simulations in this paper have been performed using the standard software routine `spm_MDP_VB_X.m`. This is freely available as Matlab code as part of academic software which can be downloaded at <http://www.fil.ion.ucl.ac.uk/spm/>.

Authors' contributions. T.P. wrote and performed the simulations using an inversion scheme written by K.J.F. T.P. and K.J.F. contributed to the writing.

Competing interests. We declare we have no competing interests.

Funding. T.P. is supported by the Rosetrees Trust (Award Number 173346). K.J.F. is a Wellcome Principal Research Fellow (Ref: 088130/Z/09/Z).

Acknowledgement. We are grateful to three anonymous reviewers for their insightful comments and suggestions.

Endnote

¹As opposed to epistemic in the philosophical sense.

Appendix A

A.1. Belief update equations

The variational free energy for the MDP model described in the main text is

$$F = -E_Q[\ln P(\tilde{o}, \tilde{s}, \pi) - \ln Q(\tilde{s}, \pi)].$$

Here, the \sim notation denotes sequences of outcomes and hidden states over time. Treating policy selection as a process of Bayesian model selection, we can consider the free energy of each policy

$$F(\pi) = -E_Q[\ln P(\tilde{o}, \tilde{s}|\pi) - \ln Q(\tilde{s}|\pi)].$$

Using the mean-field approximation $Q(\tilde{s}|\pi) = \prod_{\tau} Q(s_{\tau}|\pi)$, and factorizing the joint distribution using the conditional independencies in the generative model, this can be written as:

$$F(\pi) = -E_Q \left[\ln P(s_1) + \sum_{\tau} (\ln P(o_{\tau}|s_{\tau}) + \ln P(s_{\tau+1}|s_{\tau}, \pi) - \ln Q(s_{\tau}|\pi)) \right].$$

Taking the variational derivative of this with respect to beliefs about hidden states at a particular time we have,

$$\begin{aligned} & \frac{\delta F(\pi)}{\delta Q(s_\tau|\pi)} \\ &= -E_{\prod_{t \neq \tau} Q(s_t|\pi)} \left[\ln P(s_1) + \sum_{\tau} (\ln P(o_\tau|s_\tau) + \ln P(s_{\tau+1}|s_\tau, \pi)) \right] \\ &+ \ln Q(s_\tau|\pi). \end{aligned}$$

Setting the variational derivative to zero, the variational solution is

$$\begin{aligned} \ln Q^*(s_\tau|\pi) &= \ln P(o_\tau|s_\tau) + E_{Q(s_{\tau-1}|\pi)} [\ln P(s_\tau|s_{\tau-1}, \pi)] \\ &+ E_{Q(s_{\tau+1}|\pi)} [\ln P(s_{\tau+1}|s_\tau, \pi)] - \ln Z. \end{aligned}$$

Where Z is a partition function (i.e. a constant). In terms of the model parameters and expectations

$$\begin{aligned} \ln \varepsilon_{\pi, \tau}^* &= \boldsymbol{\zeta} \ln \mathbf{A} \cdot o_\tau + \boldsymbol{\omega} \ln \mathbf{B}_{\pi, \tau-1} \cdot s_{\pi, \tau-1} + \boldsymbol{\omega} \ln \mathbf{B}_{\pi, \tau} \cdot s_{\pi, \tau+1} \\ &- \ln Z. \end{aligned}$$

To express this as a gradient ascent, we define an error term as the difference between the current belief and the variational solution, and change the current belief to minimize this error:

$$\begin{aligned} \boldsymbol{\varepsilon}_{\pi, \tau} &= \ln \varepsilon_{\pi, \tau}^* - \ln s_{\pi, \tau} \\ \text{and } \dot{\boldsymbol{\nu}}_{\pi, \tau} &= \boldsymbol{\varepsilon}_{\pi, \tau}; \quad s_{\pi, \tau} = \sigma(\boldsymbol{\nu}_{\pi, \tau}). \end{aligned}$$

To determine the form of $\boldsymbol{\varepsilon}_{\pi, \tau}$ we have specified the form of the distributions contributing to the variational solution. In the above, the log transition probability takes the form $\boldsymbol{\omega} \ln \mathbf{B}_{\pi, \tau}$, while the log likelihood is $(\ln \mathbf{A})\boldsymbol{\zeta}$, where

$$\boldsymbol{\zeta} = \begin{bmatrix} \zeta^{(1)} & & & \\ & \zeta^{(2)} & & \\ & & \ddots & \\ & & & \ddots \end{bmatrix}.$$

There are partition functions associated with these distributions; however, these are constant with respect to the belief about the current state. This means they can be omitted in the derivation of the Bayesian belief updates for hidden states. Substituting the log probability distributions into the above, the error can be written,

$$\begin{aligned} \boldsymbol{\varepsilon}_{\pi, \tau} &= \boldsymbol{\zeta} \ln \mathbf{A} \cdot o_\tau + \boldsymbol{\omega} \ln \mathbf{B}_{\pi, \tau-1} \cdot s_{\pi, \tau-1} + \boldsymbol{\omega} \ln \mathbf{B}_{\pi, \tau} \cdot s_{\pi, \tau+1} \\ &- \ln s_{\pi, \tau}. \end{aligned}$$

This is the form used for belief updating in figure 1.

A.2. Derivation of Bayesian updates for precision parameters

In deriving belief updates for precision parameters, the partition functions cannot be omitted, as they are functions of the parameters being updated. We assume that the precision parameters are distributed according to gamma distributions, and follow a similar line of reasoning to that used to derive updates for policy precisions in previous papers. The prior distribution over the precision parameters is then:

$$P(\zeta^{(i)}) \propto \beta_z e^{-\beta_z \zeta^{(i)}}, \quad P(\omega) \propto \beta_w e^{-\beta_w \omega}.$$

The approximate posterior distributions have the same (gamma distribution) form and we will use a bold β hyperparameter to distinguish between the sufficient statistics of the posterior and prior above. A useful property of the gamma distribution, when parametrized in this way, is the following

$$\boldsymbol{\zeta} = E_{Q(\zeta)}[\boldsymbol{\zeta}] = \boldsymbol{\beta}_z^{-1}, \quad \boldsymbol{\omega} = E_{Q(\omega)}[\omega] = \boldsymbol{\beta}_w^{-1}.$$

Having defined these distributions, we can write the variational free energy

$$\begin{aligned} F &= F(\pi) + D_{KL}[Q(\pi)||P(\pi)] + D_{KL}[Q(\omega)||P(\omega)] \\ &+ \sum_i D_{KL}[Q(\zeta^{(i)})||P(\zeta^{(i)})]. \end{aligned}$$

Which can be expressed in terms of sufficient statistics (omitting constants),

$$\begin{aligned} F &= -\boldsymbol{\pi} \cdot \sum_{\tau} s_{\pi, \tau} \cdot (\ln \bar{\mathbf{A}} \cdot o_\tau + \ln \bar{\mathbf{B}}_{\pi, \tau} \cdot s_{\pi, \tau+1}) + \ln \boldsymbol{\beta}_w - \ln \beta_w \\ &+ \boldsymbol{\omega} \beta_w + \ln \boldsymbol{\beta}_z - \ln \beta_z + \boldsymbol{\zeta}^{(i)} \beta_z, \end{aligned}$$

where

$$\bar{\mathbf{A}}_{ij} = \frac{\mathbf{A}_{ij}^{\zeta^{(i)}}}{\sum_k \mathbf{A}_{kj}^{\zeta^{(i)}}} \quad \text{and} \quad \bar{\mathbf{B}}_{\pi, \tau, ij} = \frac{\mathbf{B}_{ij}^{\omega}}{\sum_k \mathbf{B}_{kj}^{\omega}},$$

correspond to normalized probability matrices. Taking the partial derivative with respect to the expected precision of state transitions gives:

$$\frac{\partial F}{\partial \boldsymbol{\omega}} = -\boldsymbol{\pi} \cdot \sum_{\tau} s_{\pi, \tau} \cdot \partial_{\boldsymbol{\omega}} \ln \bar{\mathbf{B}}_{\pi, \tau} \cdot s_{\pi, \tau+1} + \beta_w - \boldsymbol{\beta}_w.$$

Solving for zero gives the volatility expectations:

$$\boldsymbol{\beta}_w = \beta_w - \boldsymbol{\pi} \cdot \sum_{\tau} s_{\pi, \tau} \cdot \partial_{\boldsymbol{\omega}} \ln \bar{\mathbf{B}}_{\pi, \tau} \cdot s_{\pi, \tau+1}.$$

The updates for the sensory precision are obtained in the same way:

$$\boldsymbol{\beta}_z = \beta_z - \sum_{\tau} s_{\tau} \cdot \partial_{\boldsymbol{\zeta}^{(i)}} \ln \bar{\mathbf{A}} \cdot o_\tau.$$

Expressing these updates as biologically plausible gradient ascents, the resulting equations are

$$\dot{\boldsymbol{\beta}}_z = \boldsymbol{\varepsilon}_z^{(i)}; \quad \boldsymbol{\zeta}^{(i)} = \boldsymbol{\beta}_z^{-1}$$

and

$$\dot{\boldsymbol{\beta}}_w = \boldsymbol{\varepsilon}_w; \quad \boldsymbol{\omega} = \boldsymbol{\beta}_w^{-1},$$

where the errors are

$$\boldsymbol{\varepsilon}_z^{(i)} = \beta_z - \sum_{\tau} s_{\tau} \cdot \partial_{\boldsymbol{\zeta}^{(i)}} \ln \bar{\mathbf{A}} \cdot o_\tau - \boldsymbol{\beta}_z$$

$$\text{and} \quad \boldsymbol{\varepsilon}_w = \beta_w - \boldsymbol{\pi} \cdot \sum_{\tau} s_{\pi, \tau} \cdot \partial_{\boldsymbol{\omega}} \ln \bar{\mathbf{B}}_{\pi, \tau} \cdot s_{\pi, \tau+1} - \boldsymbol{\beta}_w.$$

The partial derivatives here can be expressed explicitly in terms of weighted prediction errors

$$s_{\tau} \cdot \partial_{\boldsymbol{\zeta}^{(i)}} \ln \bar{\mathbf{A}} \cdot o_\tau = \ln \mathbf{A}_{\cdot, i} \cdot \boldsymbol{\varepsilon}_{o, \tau}^{(i)}; \quad \boldsymbol{\varepsilon}_{o, \tau}^{(i)} = (o_\tau - \bar{\mathbf{A}}_{\cdot, i}) s_{\tau, i}$$

$$\begin{aligned} \text{and} \quad s_{\pi, \tau} \cdot \partial_{\boldsymbol{\omega}} \ln \bar{\mathbf{B}}_{\pi, \tau} \cdot s_{\pi, \tau+1} &= \sum_i \ln \mathbf{B}_{\pi, \tau, i} \cdot \boldsymbol{\varepsilon}_{s, \tau}^{(i)}; \quad \boldsymbol{\varepsilon}_{s, \tau}^{(i)} \\ &= (s_{\pi, \tau+1} - \bar{\mathbf{B}}_{\pi, \tau, i}) s_{\pi, \tau, i}. \end{aligned}$$

These are equivalent to the forms in figure 6, which are expressed to emphasize the plausibility of their neuronal implementation.

1. Tchaptchet A, Jin W, Braun HA. 2016 Diversity and noise in neurodynamics across different functional levels. In *Advances in Cognitive Neurodynamics (V): Proceedings of the Fifth International Conference on Cognitive Neurodynamics—2015* (eds R Wang, X Pan), pp. 681–687. Singapore, Singapore: Springer.
2. Feldman H, Friston K. 2010 Attention, uncertainty, and free-energy. *Front. Hum. Neurosci.* **4**, 215. (doi:10.3389/fnhum.2010.00215)
3. Gregory RL. 1980 Perceptions as hypotheses. *Phil. Trans. R. Soc. Lond. B* **290**, 181–197. (doi:10.1098/rstb.1980.0090)
4. Helmholtz HV. 1962 On the perceptions in general. In *Treatise on physiological optics*, vol. 3 (ed. JPC Southall), pp. 1–37. New York, NY: Dover Publications.
5. Friston K, Kilner J, Harrison L. 2006 A free energy principle for the brain. *J. Physiol. Paris* **100**, 70–87. (doi:10.1016/j.jphysparis.2006.10.001)
6. Beal MJ. 2003 Variational algorithms for approximate Bayesian inference. PhD thesis. University of London, UK.
7. Friston K. 2003 Learning and inference in the brain. *Neural Netw.* **16**, 1325–1352. (doi:10.1016/j.neunet.2003.06.005)
8. Doya K. 2008 Modulators of decision making. *Nat. Neurosci.* **11**, 410–416. (doi:10.1038/nn2077)
9. Friston K, Thornton C, Clark A. 2012 Free-energy minimization and the dark-room problem. *Front. Psychol.* **3**, 130. (doi:10.3389/fpsyg.2012.00130)
10. Friston K, Adams RA, Perrinet L, Breakspear M. 2012 Perceptions as hypotheses: saccades as experiments. *Front. Psychol.* **3**, 151. (doi:10.3389/fpsyg.2012.00151)
11. Friston K, FitzGerald T, Rigoli F, Schwartenbeck P, O’Doherty J, Pezzulo G. 2016 Active inference and learning. *Neurosci. Biobehav. Rev.* **68**, 862–879. (doi:10.1016/j.neubiorev.2016.06.022)
12. Friston K, FitzGerald T, Rigoli F, Schwartenbeck P, Pezzulo G. 2016 Active inference: a process theory. *Neural Comput.* **29**, 1–49. (doi:10.1162/NECO_a_00912)
13. Friston K, Rigoli F, Ognibene D, Mathys C, FitzGerald T, Pezzulo G. 2015 Active inference and epistemic value. *Cogn. Neurosci.* **6**, 187–214. (doi:10.1080/17588928.2015.1020053)
14. Friston K, Schwartenbeck P, FitzGerald T, Moutoussis M, Behrens T, Dolan R. 2013 The anatomy of choice: active inference and agency. *Front. Hum. Neurosci.* **7**, 598. (doi:10.3389/fnhum.2013.00598)
15. Friston K, Schwartenbeck P, FitzGerald T, Moutoussis M, Behrens T, Dolan RJ. 2014 The anatomy of choice: dopamine and decision-making. *Phil. Trans. R. Soc. B* **369**, 20130481. (doi:10.1098/rstb.2013.0481)
16. Friston KJ, Rosch R, Parr T, Price C, Bowman H. 2017 Deep temporal models and active inference. *Neurosci. Biobehav. Rev.* **77**, 388–402. (doi:10.1016/j.neubiorev.2017.04.009)
17. Mirza MB, Adams RA, Mathys CD, Friston KJ. 2016 Scene construction, visual foraging, and active inference. *Front. Comput. Neurosci.* **10**, 56. (doi:10.3389/fncom.2016.00056)
18. Bellman R. 1957 A Markovian decision process. AD0606367. Santa Monica, CA: RAND Corp. See <http://www.dtic.mil/get-tr-doc/pdf?AD=AD0606367>.
19. Carpenter GA, Grossberg S. 1988 The ART of adaptive pattern recognition by a self-organizing neural network. *Computer* **21**, 77–88. (doi:10.1109/2.33)
20. Dayan P, Hinton GE, Neal RM, Zemel RS. 1995 The helmholtz machine. *Neural Comput.* **7**, 889–904. (doi:10.1162/neco.1995.7.5.889)
21. Friston KJ, Parr T, Vries Bd. 2017 The graphical brain: belief propagation and active inference. *Netw. Neurosci.* **0**, 1–78. (doi:10.1162/NETN_a_00018)
22. Sutton RS, Barto AG. 1998 *Reinforcement learning: an introduction*. Cambridge, MA: MIT press.
23. Gibson JJ. 1966 *The senses considered as perceptual systems*. Oxford, UK: Houghton Mifflin.
24. Gibson JJ. 2014 *The ecological approach to visual perception: classic edition*. New York, NY: Psychology Press.
25. Demirdjian D, Taycher L, Shakhnarovich G, Grauman K, Darrell T. 2005 Avoiding the ‘streetlight effect’: tracking by exploring likelihood modes. In *Tenth IEEE International Conference on Computer Vision (ICCV’05), Beijing, China, 17–21 October*, vol. 1, pp. 357–364. Piscataway, NJ: IEEE.
26. Fillmore MT, Rush CR, Abrams BD. 2005 d-Amphetamine-induced enhancement of inhibitory mechanisms involved in visual search. *Exp. Clin. Psychopharmacol.* **13**, 200. (doi:10.1037/1064-1297.13.3.200)
27. Mathôt S, Dalmajer E, Grainger J, Van der Stigchel S. 2014 The pupillary light response reflects exogenous attention and inhibition of return. *J. Vis.* **14**, 7. (doi:10.1167/14.14.7)
28. Gabay S, Pertzov Y, Henik A. 2011 Orienting of attention, pupil size, and the norepinephrine system. *Atten. Percept. Psychophys.* **73**, 123–129. (doi:10.3758/s13414-010-0015-4)
29. Pertzov Y, Bays PM, Joseph S, Husain M. 2013 Rapid forgetting prevented by retrospective attention cues. *J. Exp. Psychol. Hum. Percept. Perform.* **39**, 1224–1231. (doi:10.1037/a0030947)
30. Burak Y, Fiete IR. 2012 Fundamental limits on persistent activity in networks of noisy neurons. *Proc. Natl Acad. Sci. USA* **109**, 17 645–17 650. (doi:10.1073/pnas.1117386109)
31. Posner MI, Rafal RD, Choate LS, Vaughan J. 1985 Inhibition of return: neural basis and function. *Cogn. Neuropsychol.* **2**, 211–228. (doi:10.1080/02643298508252866)
32. Klein RM. 2000 Inhibition of return. *Trends Cogn. Sci.* **4**, 138–147. (doi:10.1016/S1364-6613(00)01452-2)
33. Yu AJ, Dayan P. 2005 Uncertainty, neuromodulation, and attention. *Neuron* **46**, 681–692. (doi:10.1016/j.neuron.2005.04.026)
34. Dayan P, Yu AJ. 2001 ACh, uncertainty, and cortical inference. In *Advances in Neural Information Processing Systems 14 (NIPS 2001)*, (eds TG Dietterich, S Becker, Z Ghahraman), pp. 189–196. Neural Information Processing Systems Foundation, Inc.
35. Yu AJ, Dayan P. 2002 Acetylcholine in cortical inference. *Neural Netw.* **15**, 719–730. (doi:10.1016/S0893-6080(02)00058-8)
36. Dayan P, Yu AJ. 2006 Phasic norepinephrine: a neural interrupt signal for unexpected events. *Netw. Comput. Neural Syst.* **17**, 335–350. (doi:10.1080/09548980601004024)
37. Lavine N, Reuben M, Clarke P. 1997 A population of nicotinic receptors is associated with thalamocortical afferents in the adult rat: laminar and areal analysis. *J. Comp. Neurol.* **380**, 175–190. (doi:10.1002/(SICI)1096-9861(19970407)380:2<175::AID-CNE3>3.0.CO;2-0)
38. Sahin M, Bowen WD, Donoghue JP. 1992 Location of nicotinic and muscarinic cholinergic and μ -opioid receptors in rat cerebral neocortex: evidence from thalamic and cortical lesions. *Brain Res.* **579**, 135–147. (doi:10.1016/0006-8993(92)90752-U)
39. Shipp S. 2007 Structure and function of the cerebral cortex. *Curr. Biol.* **17**, R443–R449. (doi:10.1016/j.cub.2007.03.044)
40. Disney AA, Aoki C, Hawken MJ. 2007 Gain modulation by nicotine in macaque V1. *Neuron* **56**, 701–713. (doi:10.1016/j.neuron.2007.09.034)
41. Gil Z, Connors BW, Amitai Y. 1997 Differential regulation of neocortical synapses by neuromodulators and activity. *Neuron* **19**, 679–686. (doi:10.1016/S0896-6273(00)80380-3)
42. Arroyo S, Bennett C, Hestrin S. 2014 Nicotinic modulation of cortical circuits. *Front. Neural Circuits* **8**, 30. (doi:10.3389/fncir.2014.00030)
43. Eckenstein FP, Baughman RW, Quinn J. 1988 An anatomical study of cholinergic innervation in rat cerebral cortex. *Neuroscience* **25**, 457–474. (doi:10.1016/0306-4522(88)90251-5)
44. Herrero JL, Roberts MJ, Delicato LS, Gieselmann MA, Dayan P, Thiele A. 2008 Acetylcholine contributes through muscarinic receptors to attentional modulation in V1. *Nature* **454**, 1110–1114. (doi:http://www.nature.com/nature/journal/v454/n7208/supinfo/nature07141_s1.html)
45. Aston-Jones G, Cohen JD. 2005 An integrative theory of locus coeruleus–norepinephrine function: adaptive gain and optimal performance. *Annu. Rev. Neurosci.* **28**, 403–450. (doi:10.1146/annurev.neuro.28.061604.135709)
46. Koss MC. 1986 Pupillary dilation as an index of central nervous system α 2-adrenoceptor activation. *J. Pharmacol. Methods* **15**, 1–19. (doi:10.1016/0160-5402(86)90002-1)

47. Kahneman D, Beatty J. 1966 Pupil diameter and load on memory. *Science* **154**, 1583. (doi:10.1126/science.154.3756.1583)
48. Mason ST, Fibiger HC. 1977 Altered exploratory behaviour after 6-OHDA lesion to the dorsal noradrenergic bundle. *Nature* **269**, 704–705. (doi:10.1038/269704a0)
49. Wendlandt S, File SE. 1979 Behavioral effects of lesions of the locus ceruleus noradrenaline system combined with adrenalectomy. *Behav. Neural Biol.* **26**, 189–201. (doi:10.1016/S0163-1047(79)92579-2)
50. Tait DS, Brown VJ, Farovik A, Theobald DE, Dalley JW, Robbins TW. 2007 Lesions of the dorsal noradrenergic bundle impair attentional set-shifting in the rat. *Eur. J. Neurosci.* **25**, 3719–3724. (doi:10.1111/j.1460-9568.2007.05612.x)
51. Fibiger HC, Mason ST. 1978 The effects of dorsal bundle injections of 6-hydroxydopamine on avoidance responding in rats. *Br. J. Pharmacol.* **64**, 601–605. (doi:10.1111/j.1476-5381.1978.tb17322.x)
52. Berridge CW, Waterhouse BD. 2003 The locus coeruleus–noradrenergic system: modulation of behavioral state and state-dependent cognitive processes. *Brain Res. Rev.* **42**, 33–84. (doi:10.1016/S0165-0173(03)00143-7)
53. Arnsten AFT, Goldman-Rakic PS. 1984 Selective prefrontal cortical projections to the region of the locus coeruleus and raphe nuclei in the rhesus monkey. *Brain Res.* **306**, 9–18. (doi:10.1016/0006-8993(84)90351-2)
54. Sara SJ, Hervé-Minvielle A. 1995 Inhibitory influence of frontal cortex on locus coeruleus neurons. *Proc. Natl Acad. Sci. USA* **92**, 6032–6036. (doi:10.1073/pnas.92.13.6032)
55. Schwartenbeck P, Friston K. 2016 Computational phenotyping in psychiatry: a worked example. *eNeuro* **3**, ENEURO.0049-16.2016. (doi:10.1523/ENEURO.0049-16.2016)
56. Zaborszky L, Gaykema RP, Swanson DJ, Cullinan WE. 1997 Cortical input to the basal forebrain. *Neuroscience* **79**, 1051–1078. (doi:10.1016/S0306-4522(97)00049-3)
57. Galea JM, Bestmann S, Beigi M, Jahanshahi M, Rothwell JC. 2012 Action reprogramming in parkinson's disease: response to prediction error is modulated by levels of dopamine. *J. Neurosci.* **32**, 542. (doi:10.1523/JNEUROSCI.3621-11.2012)
58. Frank MJ. 2005 Dynamic dopamine modulation in the basal ganglia: a neurocomputational account of cognitive deficits in medicated and nonmedicated parkinsonism. *J. Cogn. Neurosci.* **17**, 51–72. (doi:10.1162/0898929052880093)
59. Goldman-Rakic PS, Castner SA, Svensson TH, Siever LJ, Williams GV. 2004 Targeting the dopamine D1 receptor in schizophrenia: insights for cognitive dysfunction. *Psychopharmacology (Berl)* **174**, 3–16. (doi:10.1007/s00213-004-1793-y)
60. Adams RA, Stephan KE, Brown HR, Frith CD, Friston KJ. 2013 The computational anatomy of psychosis. *Front. Psychiatry* **4**, 47. (doi:10.3389/fpsy.2013.00047)
61. Howes OD, Kapur S. 2009 The dopamine hypothesis of schizophrenia: version III—the final common pathway. *Schizophr. Bull.* **35**, 549–562. (doi:10.1093/schbul/sbp006)
62. Lombardo S, Maskos U. 2015 Role of the nicotinic acetylcholine receptor in Alzheimer's disease pathology and treatment. *Neuropharmacology* **96**, Part B, 255–262. (doi:10.1016/j.neuropharm.2014.11.018)
63. Whitehouse PJ, Price DL, Clark AW, Coyle JT, DeLong MR. 1981 Alzheimer disease: evidence for selective loss of cholinergic neurons in the nucleus basalis. *Ann. Neurol.* **10**, 122–126. (doi:10.1002/ana.410100203)
64. Blier P, El Mansari M. 2007 The importance of serotonin and noradrenaline in anxiety. *Int. J. Psychiatry Clin. Pract.* **11**, 16–23. (doi:10.1080/13651500701388310)
65. Gören MZ, Cabadak H. 2014 Noradrenaline and post-traumatic stress disorder. In *Comprehensive guide to post-traumatic stress disorder* (eds CR Martin, VR Preedy, VB Patel), pp. 1–16. Cham, Switzerland: Springer International Publishing.
66. Moret C, Brolley M. 2011 The importance of norepinephrine in depression. *Neuropsychiatr. Dis. Treat.* **7**, 9–13. (doi:10.2147/NDT.S19619)
67. Halliday G, Cullen K, Harding A. 1993 Neuropathological correlates of memory dysfunction in the Wernicke–Korsakoff syndrome. *Alcohol. Alcohol.* **2**, 245–251.
68. Mair R, Anderson CD, Langlais P, McEntee W. 1985 Thiamine deficiency depletes cortical norepinephrine and impairs learning processes in the rat. *Brain Res.* **360**, 273–284. (doi:10.1016/0006-8993(85)91243-0)
69. Oke A, Keller R, Mefford I, Adams RN. 1978 Lateralization of norepinephrine in human thalamus. *Science* **200**, 1411. (doi:10.1126/science.663623)
70. Kim M, Barrett AM, Heilman KM. 1998 Lateral asymmetries of pupillary responses. *Cortex* **34**, 753–762. (doi:10.1016/S0010-9452(08)70778-0)
71. Malhotra PA, Parton AD, Greenwood R, Husain M. 2006 Noradrenergic modulation of space exploration in visual neglect. *Ann. Neurol.* **59**, 186–190. (doi:10.1002/ana.20701)
72. Daw ND, Doya K. 2006 The computational neurobiology of learning and reward. *Curr. Opin Neurobiol.* **16**, 199–204. (doi:10.1016/j.conb.2006.03.006)
73. Doya K. 2002 Metalearning and neuromodulation. *Neural Netw.* **15**, 495–506. (doi:10.1016/S0893-6080(02)00044-8)