Adaptive learning is structure learning in time

Linda Q. Yu^{1*}, Robert C. Wilson² & Matthew R. Nassar¹

¹Carney Institute for Brain Sciences, Brown University, 164 Angell Street, Providence, RI, 02912, USA

² Department of Psychology, University of Arizona, Tucson, AZ, 85721, USA

*Corresponding author:

linda_yu@brown.edu

Abstract

People use information flexibly. They often combine multiple sources of relevant information over time in order to inform decisions with little or no interference from intervening irrelevant sources. They adjust the degree to which they use new information over time rationally in accordance with environmental statistics and their own uncertainty. They can even use information gained in one situation to solve a problem in a very different one. Learning flexibly rests on the ability to infer the context at a given time, and therefore knowing which pieces of information to combine and which to separate. We review the psychological and neural mechanisms behind adaptive learning and structure learning to outline how people pool together relevant information, demarcate contexts, prevent interference between information collected in different contexts, and transfer information from one context to another. By examining all of these processes through the lens of optimal inference we bridge concepts from multiple fields to provide a unified multi-system view of how the brain exploits structure in time to optimize learning.

Keywords: structure learning, adaptive learning, grid code, Bayesian inference, reversal learning, event segmentation

Introduction

On June 5, 2020, the jobs numbers in the US for May came out. The previous day, the media had speculated that the unemployment rate for May would only go up from April's historic highs of 14.7% -- the question was only how much higher. Unexpectedly, the unemployment rate for May turned out to be 13.3%, a decrease. The stock market jumped in response. Was this a sign that the historic downturn in March caused by the coronavirus was a mere blip and the economy would rapidly bounce back? Or was it a mirage that will dissipate into a prolonged recession? The dip in unemployment numbers from April to May was, in reality, fairly small, but the prediction error from the widely expected increase to the reality of a decrease was much larger. Investors evidently took this positive error as a sign of the impending economic recovery and acted accordingly, but were they correct to do so?

Our brain generates expectations all the time – influential theories have even argued this is its main purpose (Friston, 2012). When expectations are violated, we must decide if we are now in a new regime and shift our behavior drastically, or dismiss this new observation as aberrant and continue with our previous policies. In comparison to the above scenario, there was a similar “rally” in the stock market in 2008 after the initial crash, but that bounce disappeared soon after as the stock market continued its fall and the nation was plunged into the Great Recession. Two situations that look superficially similar can be caused by completely different realities. How much we adjust our behavior in response to a surprising new observation thus must depend on our inference of what context we are in. If an investor thinks she is in a situation much like the Great Recession, she might not put more money into a stock market that will eventually plunge lower, but if she think she is in a true recovery, then she would want to put her money in now to benefit from the new gains. The process of adjusting how much incoming information affects our beliefs is known as *adaptive learning*, while the

process of inferring the underlying context we refer to as *structure learning* (Box 1). As this example illustrates, the two processes must be intimately related in order to produce adaptive behaviors.

Here we propose that adaptive learning and structure learning ultimately address the same underlying question: which pieces of information do we want to combine, and which do we want to keep separate in order to make predictions about the future? For example, do we think the lessons from the 2008 recession should apply to our current investment choices, or not? To solve this problem, we propose that people make use of hidden (or latent) “states” that capture the set of information relevant to a given action or decision (Wilson, Takahashi, Schoenbaum, & Niv, 2014). These states serve three key functions. First, they facilitate prediction, when a state is active previous information assigned to that state is available (e.g. if we are in a recession we should expect things to get worse). Second, they allow us to infer when a context has changed; when predictions of the currently active state are violated by surprising information, we may activate a different state or create an entirely a new state (e.g. positive employment numbers may indicate a recession is ending). Third, the states allow us to separate associations and information between states, so that they do not interfere with each other (e.g., we are no longer in a period of continued growth prior to the pandemic). In this article, we will review the literatures on adaptive learning and structure learning, using this state perspective as a mechanism for how the brain uses prediction errors to form new states, segregate independent sources of information across states, and transfer more global structure information between states.

Figure 1. The stock market example illustrating the problem of learning in changing environments. When faced with a large prediction error (the coronavirus crisis reaching the US), information should be segregated from the previous state (i.e., no longer in period of continuous growth as before). When a second prediction error happens (better than expected jobs number in June 2020), investors must decide whether they should believe in continued recovery, or if they should transfer their mental model from the 2008 recession. Stock market images are of the Standard and Poor (S&P) index of 500 US common stock in 2020 (left) and 2008 (right).

Transitions Between States

States Facilitate Prediction

When we observe something surprising, such as better than expected employment numbers in June, a key question that we are likely to ask ourselves is, what will happen next? If we can make good predictions about the future, we will be able to respond better to the change. However, the predictions that we make depend critically on what we believe about the causes of the change. Perhaps the decreased unemployment rate was caused by the end of the pandemic and everything going back to normal. Perhaps people went back to work because their one-time stimulus check ran out and they needed the money. Perhaps it was just a random positive blip on the way to a much bigger crash. To make the best predictions we need to consider many different causes like these and weigh them according to how likely we believe them to be true (Box 2).

One way to formulate these different causal explanations mathematically is by defining *latent states*, which exert a causal influence on our observations, but cannot be observed directly (Gershman & Niv, 2010). Latent states can facilitate predictions by grouping observations together when they appear to share a common cause. Written mathematically, given the state S_t at time t , and the previous data points associated with this state, X_{S_t} , we can make predictions about future data X_{t+1} , according to the probability

$$P(X_{t+1}|S_t, X_{S_t})$$

The dependence on past data associated with this state, X_{S_t} , is critical here. This implies that the distribution over predicted outcomes is shaped by these past outcomes such that, if we are in the correct state, our predictions are likely to be good. Of course, if our predictions turn out to be bad, then we may need to change the currently activated state.

Surprising Observations Can Indicate Changes in Underlying State

If the state is truly latent (i.e. it is not directly observed) and/or if the state can change unexpectedly over time, then we can never be completely sure which state we are in. We therefore need to *infer* the currently active state based on the observed data. Formally, a probability distribution over the currently active state can be computed using Bayes' rule:

$$P(S_t|X_t, S_{1:t-1}, X_{1:t-1}) = \frac{P(X_t|S_{1:t}, X_{1:t-1})P(S_t|S_{1:t-1}, X_{1:t-1})}{P(X_t|S_{1:t-1}, X_{1:t-1})}$$

where $P(S_t|X_t, S_{1:t-1}, X_{1:t-1})$ is the distribution over the state given the previous observations and previously active states, $P(X_t|S_{1:t}, X_{1:t-1})$ is the likelihood of the observed data given each

state, $P(S_t|S_{1:t-1}, X_{1:t-1})$ is the prior transition probability of switching from one state to another, and $P(X_t|S_{1:t-1}, X_{1:t-1})$ is a normalizing constant.

The idea of surprise flows naturally from this equation, as it is being observations for which the likelihood of the data given the state $P(X_t|S_{1:t}, X_{1:t-1})$ is very low. At the same time that we are assessing the likelihood of the data under a state, we are also updating our beliefs about the parameters that describe the observations that are likely in that state (eg., mean, variance, etc.) using the data we have assigned to it (Box 3).

The manner in which a surprising event should prompt us to revise our mental model (i.e., belief updating) depends on the perceived source of that surprising event. For example, the prediction error could be generated from noise (like the minute-to-minute fluctuations in daily stock prices), in which case it does not reflect any change in the source of the observations and should not lead us to fundamentally revise our beliefs. We should update our beliefs more drastically, however, if the change is believed to be informative, or reflecting an underlying change in the state (e.g., the price of the stock of a national department store chain after the stay-at-home orders went into effect). But sometimes, a dramatic change happens that is still not reflective of an underlying change (as in the 2010 Flash Crash, where a rapid plunge in the stock market lasted 36 minutes). In psychology, such events are referred to as “oddballs”, and may look like changepoints, but do not persist in time and thus should not affect predictions about future observations, and therefore, behavior.

Even in environments where changes tend to persist, more subtle differences in statistical context impact behavior. People’s beliefs about the nature of changes additionally depend on whether the overall context is volatile (changes in state occur often) or stable (changes in state do not happen very often). They generally adapt more in volatile (rapidly changing) environments in which beliefs based on previous experience can become rapidly outdated and irrelevant to the problem of predicting future outcomes (Behrens, Woolrich, Walton, & Rushworth, 2007; Browning, Behrens, Jocham, O’Reilly, & Bishop, 2015; Pulcu & Browning, 2017). So, jumps in stock prices in the days following the global coronavirus shutdown would be perceived as more meaningful due to the sparsity of data for this new context, and thus more influential, than those occurring during the period of stability prior to it. Another important aspect of statistical context that affects how people update beliefs is the degree to which previously encountered contexts are repeated – people can recognize and exploit repeated contexts to eliminate the need for relearning after a state change (Collins & Koechlin, 2012). A model that hopes to accurately predict the future must take into account all of these varied scenarios.

Exploiting Structure in a Time Series with Latent States

While the literature on adaptive learning provides a descriptive and normative account of the behavioral and neural correlates of belief updating in specific scenarios, it has so far fallen short on offering a mechanistic and plausible model of how this process arises in the brain more generally. The degree to which new experience shapes immediate future behavior is known as the learning rate, but we argue that this terminology may be misleading, and that in fact the mechanisms underlying behavioral adjustment have less to do with a quantity of learning, and more to do with the dynamic representations to which associations are formed.

Existing implementational models are limited in the range of behavior they can explain, whereas human behavior is incredibly flexible and capable of optimizing learning in many ways depending on task structure. For example, recent work has shown that people can separately and simultaneously adjust learning rates for different affective outcomes (Pulcu & Browning, 2017). Another adaptive aspect of human behavior that has proven troubling for neural models is that humans are capable of drastically altering strategy in the face of different task structures; for example, people rapidly increase learning to surprising information in changing environments, but minimize the weight to surprising information in contexts with uninformative oddballs (Nassar, Bruckner, & Frank, 2019). These forms of flexibility make it difficult to imagine the range of adjustments in learning rate observed being produced by a simple fixed architecture neural network, even if this network had access to critical factors necessary for learning rate adjustment within a single context. What governs a learning rate that dynamically adjusts according to the statistics of the environment? What makes one significant prediction error predictive of an informative state change, and another of the same magnitude an “oddball”, subject to a completely different set of neural responses and behavior? Here we consider the possibility that dynamic learning rate is a natural consequence of structure learning over time.

We will discuss the framework of dynamic structure learning in the context of three paradigms that are often used in psychology and neuroscience research, changepoint detection, oddball, and reversal learning (Figure 2A). The three tasks can each be described in terms of the same generative and inference structure described above that consists of a set of states, a parameterized function through which states generate observations, and a transition function that defines how states change from moment-to-moment. The key difference between the inference models for the three tasks is in the transition function, which defines the probability with which a given state on trial t could transition to any other state at time $t+1$. The transition function, formally the conditional probability distribution $P(S_{t+1}/S_t)$, is depicted for each task in Figure 2B. In principle, these transition functions could be used to control and update latent state representations in the brain, thereby controlling the active representation to which new observations are bound (Figure 2C). The implications of differential transition functions for behavior and neural representation are described in more detail for each task separately below.

Figure 2. A) Diagram representations of the timeline of changepoint, oddball, and reversal learning tasks, with blue dots representing outcomes on individual trials, and the orange line representing the underlying state (mean). The black boxes outline the surprising events in each case (labeled 1 and 2, respectively). The color of the shading of the background represents distinct states that are either created, or returned to upon a prediction error. B) State transition probability functions for changepoint (top), oddball (middle), and reversal (bottom) paradigms. H

denotes the hazard rate, or the probability of transitioning to a different state. Each box indicates the probability of transitioning from the row state to the column state. The ellipses in the changepoint transition function indicates an ongoing possibility of transition to new states beyond those depicted. C) A simplified representation of the neural network changes proposed in Razmi & Nassar (2020) corresponding to each task condition, occurring around the time of trial t (second surprising event) as depicted in A. The top layer denotes the output, and the bottom layer of circles represents input nodes of the network, a subset of which are activated and colored according to the current latent state. When a surprising event (e.g. trial t) happens, the latent state shifts to a new node. The lines between the layers denote the weights of the connections, with the thicker lines indicating strengthened weights. Thick gray lines indicate a pre-existing strong connection, thick yellow lines indicate a newly strengthened connection, and thick gray lines with a yellow center (as in bottom row) indicates a pre-existing strong connection that has had a new update.

Changepoint task. In a changepoint task, the underlying generative source of outcomes (e.g., the mean of a Gaussian distribution) stays the same for a period of time, then shifts abruptly to a new source (Fig 2A top row). The transition function for this task specifies that with each data point, you are likely to stay in the current state, but occasionally (i.e., a changepoint), you would transition to a new state (Fig 2B top row). As this type of task never returns to a previously experienced state, a new state is formed at each changepoint (Adams & MacKay, 2007; Wilson, Nassar, & Gold, 2010). When the probability of seeing the data given a state, $P(X_t|S_t)$, is low for all previously encountered states, we need to switch into a completely new state. Since transitions in the changepoint environment are always to a new state, the parameters governing that state must be learned from scratch. This learning from scratch leads observations occurring immediately after transitions to carry more weight, providing an intuition for higher rates of learning after changepoints and during periods of volatility. Practically, in a neural network model, this could be implemented as a shift in the neural population encoding the active latent state, which in principle could be triggered through observation of a surprising event (Razmi & Nassar, 2020) (Fig 2C top row). The neural network uses an approximation to the Bayesian inference process described above to infer when the likelihood of a state change is high, and changes its weights to a new neural population representing a novel latent state. So long as the newly activated latent state node has weak initial connections to the output layer, weight adjustments in accordance with the surprising observation would dominate the output of the network on a subsequent trial. In accordance with the transition function, the network expects to stay in that newly active node in the following trial (blue neuron in Fig 2C, top row), meaning that the recently learned weights dominate upcoming behavior (thick yellow line), which behaviorally would be interpreted as rapid learning.

In principle, the probability of a changepoint can be computed as the marginal probability that the state on trial t does not equal that on the previous state, but it is often approximated in models that collapse the many possible state transition histories (eg. mental models) down to just two: are we in the same state or a different one (Box 4)? Changepoint probability has been used widely in the adaptive learning literature to characterize participants' tendency to learn rapidly from surprising events that outcome contingencies have changed (Krishnamurthy, Nassar, Sarode, & Gold, 2017; McGuire, Nassar, Gold, & Kable, 2014; Nassar, McGuire, Ritz, & Kable, 2019; Nassar et al., 2012; Wilson, Nassar, & Gold, 2013).

Oddball task. A prediction error is not informative, even if large, if it is believed to be a statistical outlier, or “oddball”, that will not relate to future events (d’Acremont & Bossaerts, 2016; Nassar, Bruckner, et al., 2019; O’Reilly et al., 2013). In an oddball task, the state stays constant, or drifts according to a continuous process (e.g., Gaussian random walk) leading typical observations to reflect small errors in prediction. In contrast, large statistical outliers occur on occasion (Fig 2B middle row). The transition function for the oddball task suggests that you are likely to stay in the typical state (the “normal” state) with only a small chance of transitioning to the oddball state. However, once in the oddball state, you are very likely to return to the normal state on the subsequent timestep (Fig 2B middle row). Because surprisingly large errors are diagnostic of the oddball state (Box 4), they could activate a new set of neurons representing that state (Fig 2C middle row). Note that in the neural network model, both oddballs and changepoints trigger the same event at the level of neural representation: a new state formation. However, this same process has opposing effects on learning, as observed in behavior (d’Acremont & Bossaerts, 2016; Nassar, Bruckner, et al., 2019). In the changepoint case, the distinct population of neurons representing a new state serves as a “clean slate” for learning: old behavior does not need to be un-learned, as in model-free reinforcement learning, and instead new observations are integrated together. This has the effect of a speeded learning rate after a change point. In the oddball case, the new population of neurons representing the oddball serves to isolate it. This occurs through the transition function: in the oddball task, the “oddball” state is typically followed by the “normal” state (Fig 2B). Assuming that this transition occurs in the latent state representation prior to the subsequent trial, all learning that is attributed to the oddball event will be stored in the weights projecting from an inactive (green) neuron, thereby minimizing the degree to which the oddball event can affect subsequent behavior (Fig 2C). In this way, differential transition functions in changepoint and oddball tasks promote opposite learning behaviors in the face of surprising information, and structured neural state representations that obey relevant transition functions naturally capture this distinction.

Reversal learning. In reversal learning, participants go back and forth between two states, where one type of association is rewarded but not another in State 1, and vice versa in State 2 (Fig 2A bottom row). The transition function for the reversal task defines remaining in the current state to be highly probable, but in the case of a transition, specifies it to occur to the other state (thus allowing a return to a previously experienced state, unlike the changepoint task) (Fig 2B bottom row). Practically, we return to the other state when the probability of the data given the state $P(X_t|S_t)$ is low for the state we are in now, but not for a state we have previously encountered. In the neural network model, two populations of neurons represent distinct states; here, a large prediction error indicates the transition to the other state (Fig 2C bottom row), where pre-existing weights are newly updated (thick gray lines with yellow center in Fig 2C). Again, this separation allows learning to be segregated between the states and not affect each other. However, the key difference afforded by the reversal transition function is transitions to familiar states. Learning onto state representations that obey this transition structure allows information to be effectively transferred across time, for example allowing earliest data collected in the reversal panel of figure 2A to inform predictions

immediately after recognition of a transition back to state 1, thereby minimizing the need for relearning.

In the network, this occurs through retrieval of previous associations stored in the weights projecting from the state 1 neuron. This leads to explanation of interesting psychological phenomena in extinction and reversal learning. Specifically, inferences of whether the subject is in a new state or old one depending on the strength of the evidence dictates whether extinction is persistent (Gershman, Blei, & Niv, 2010; Gershman, Jones, Norman, Monfils, & Niv, 2013; Gershman, Monfils, Norman, & Niv, 2017), or whether learning after a reversal is quickly learnt after probabilistic reinforcement or overtraining (Donoso, Collins, & Koehlin, 2014; Lloyd & Leslie, 2013).

Real world scenarios will deviate from the exact transition functions that we define here. For example, while real life often repeats itself like the reversal task, we also do encounter new contexts upon occasion. The Chinese Restaurant Process (CRP) provides one way to formalize the idea of either returning to a previously known state or go to a new one: new observations are either assigned to an old state or a new one, with more “popular” states that has more observations assigned to it being more likely (Collins & Frank, 2013; Franklin, Norman, Ranganath, Zacks, & Gershman, 2020; Gershman, Radulescu, Norman, & Niv, 2014). The so-called “sticky” variant of the CRP models the tendency of events to persist in time. Yet even the sticky CRP likely underestimates the complexity of transition functions that people use in real life, where certain states may persist in time, but also tend to terminate upon entering other specific states. We hope to provide a formalism that not only links existing models in the literature, but provides a way to think beyond existing paradigms to flexible behavior observed in the wild.

Another relevant issue for models is that in the real world, we are not always certain that the state we are representing is the correct one. In such cases it is advantageous to keep alternative paths we could have taken through the state space, or counterfactuals, alive as possibilities. Representation of alternative options have been found in frontopolar cortex (Boorman, Behrens & Rushworth, 2011), and people have shown to be able to “go back” and assign outcomes to causes they only learn about later on (Moran et al., 2021). The neural network model we describe (Razmi & Nassar, 2020) indeed suffers if it only keeps track of the outcome with the highest maximum likelihood – but in cases where explicit probabilistic inference is possible, keeping track of the top few candidates improves predictions drastically (Franklin et al., 2020; Wilson et al., 2010).

There have been numerous other computational approaches to latent state inference. Particularly, Hidden Markov Model (HMM) and semi-HMM approaches allow the states to keep track of past history in a tractable way, while disambiguating states that have similar observations (Ebitz, Albarran, & Moore, 2018; Starkweather, Babayan, Uchida, & Gershman, 2017; Wierstra & Wiering, 2004). Other computational models have tackled change points, and state switching using the CRP (Collins & Koehlin, 2012; Wilson et al., 2010). However, it is difficult to explain the change point vs. oddball result with these models, because it requires knowledge, not of observations or of returns, but of the transition structure – and in the changepoint case that transition structure is conditional on previous state, violating simplifying assumptions in general models that try to learn this structure. The neural network model we

laid out here (Razmi & Nassar, 2020) is simply fed the transition structure; it's a hard problem to learn the transition structure itself, but recent work has made exciting progress in this direction (Heald, Lengyel, & Wolpert, 2020).

To summarize, our structure learning account of adaptive learning would require the following components: 1) estimation of a given latent state and its parameters allows us to predict observations; 2) a surprise signal when the observed data seems unlikely given those predictions and triggers an updating of the state, and 3) the transition function that maps the current state onto future ones (Fig 2B), as well as an representation of the active state (Fig 2A). In the next section, we will delineate how these components map onto neural processes.

Neural Mechanisms of Adaptive Learning

Here we review empirical data on the brain regions and neuromodulatory systems involved in adaptive learning with a focus on the computations laid out in the previous section. Stimulus-action values are a particularly important parameter of real-world prediction models – in many cases accurate assessment of action values could mean the difference between life and death. There has been a litany of work focused on how dopamine signed reward prediction error signals from the VTA facilitate learning of stimulus action values in striatal medium spiny cells that in turn facilitate selection of rewarding actions (Frank, Seeberger, & O'Reilly, 2004; Maes et al., 2020; Schultz, Dayan, & Montague, 1997; Waelti, Dickinson, & Schultz, 2001). While state action values sculpted by dopamine signaling may be a particularly relevant “parameter” that needs to be learned, the dopamine system has also been proposed to underlie learning of stimulus-stimulus associations, suggesting that it may shape more complex associations to a latent state, which in our framework could be thought of as learning higher order parameters (Langdon et al., 2018; Sharpe et al., 2017). In principle, other high order parameters, such as the variability over possible outcomes might also be computed, and an influential theory has suggested that they might be linked to central acetylcholine (Reimer et al., 2016; Yu & Dayan, 2005). Nonetheless, adaptive learning requires not only learning to predict the outcomes of actions within a state, but also to rapidly disengage from those predictions in the face of a changepoint, or recognize the relevance of an alternative state in the face of a reversal.

Neural mechanisms of surprise promoting state transition. One critical element necessary for this sort of disengagement is computing surprise, for purposes of detecting state transitions (Box 3). In simple tasks, the initial mismatch between expectations and observations occur in the primary sensory cortex associated with the modality of the task (Knill & Richards, 1996; McGuire et al., 2014; Meyniel & Dehaene, 2017), as well as in higher order regions including parietal, prefrontal, and anterior cingulate cortex where surprise signals are domain general (McGuire et al., 2014; Meyniel & Dehaene, 2017; O'Reilly et al., 2013). These signals are also conveyed to neuromodulatory systems, in particular the locus coeruleus/norepinephrine system, leading to higher levels of noradrenergic signaling after surprising stimuli across domains (Aston-Jones & Bloom, 1981; Foote, Aston-Jones, & Bloom, 1980). The LC/NE system provides widespread neuromodulation across cortex, positioning it to translate a surprising

stimulus in into a broader reorganization of cortical activity (Bouret & Sara, 2005), that could potentially change the active latent state representation.

Activity of the LC/NE system can be measured by proxy through transient pupil dilations and, more speculatively, the P300 event related potential (Filipowicz, Glaze, Kable, & Gold, 2020; Jepma et al., 2016; Joshi, Li, Kalwani, & Gold, 2016; Krishnamurthy et al., 2017; Nassar et al., 2012; O'Reilly et al., 2013; Reimer et al., 2016; Vazey, Moorman, & Aston-Jones, 2018). The extent to which people abandon their priors and update their beliefs in response to a changepoint is reflected in the size of the phasic pupil dilation (Krishnamurthy et al., 2017; Nassar et al., 2012), the magnitude of the P300 response (Fischer & Ullsperger, 2013; Jepma et al., 2016), and the activation of the frontoparietal network, notably including the dorsal anterior cingulate and posterior parietal cortex (Behrens et al., 2007; d'Acremont & Bossaerts, 2016; Kao et al., 2020; McGuire et al., 2014; Meyniel & Dehaene, 2017; Nour et al., 2018; O'Reilly et al., 2013). While these pupil and EEG findings have been interpreted in terms of NE exerting a direct influence on learning (Yu & Dayan, 2005), a broader look at the data suggests that this might not be the case. P300 predicts more learning in the context of changepoints, but indeed predicts less learning in oddball environments (Nassar, Bruckner, et al., 2019). Furthermore, transient pupil dilations to oddball events can persist longer than those to changepoints, such that across both contexts late pupil dilations have a negative relationship to learning (O'Reilly et al., 2013).

The neural network model described above might rectify some of the counterintuitive findings with respect to how brain signals relate to learning in different statistical contexts. We propose that NE signals the need for a state transition, rather than the need to learn per se. Through this lens, consider the finding that a larger feedback-locked P300 magnitude corresponds to more learning from the feedback that accompanies it in changing contexts, but corresponds to less learning in an oddball context. Considering the P300 as a “state transition” signal provides a reasonable way to think about this result. The P300 reflects more learning in changepoint conditions when a state transition is necessary to reduce the impact of previous observations (Fig 2, changepoint), but less learning in an oddball context where state transitions are used to minimize the interference associated with outlying datapoints (Fig 2 oddball; Nassar, Bruckner & Frank, 2019). When NE is measured by proxy through pupil diameter, which is slow and involves considerable temporal smearing, the resultant signal from an oddball event might combine two transitions (to the oddball state and back again) and thus be larger than for a changepoint, potentially explaining the discrepancy between previous studies that have linked pupil diameter either positively or negatively to learning (Nassar et al., 2012; O'Reilly et al., 2013). Furthermore, different learning rates for different affective outcomes may be possible through using a multi-dimensional state space (Pulcu and Browning, 2018). Positive and negative outcomes that vary independently could be represented by separate latent features, that then could be combined to form a decision. Mechanistically, the different features could be coded by different but overlapping neuronal populations, much like as has been proposed for dopamine prediction errors (Langdon et al., 2018). In this way, latent state might be constructed compositionally, thereby allowing transfer of some types of knowledge across contexts and preventing interference that would be caused by transfer of others.

This view of the role of NE in latent state transition is compatible with ideas about the longer-term role of NE in driving sensory gain in uncertain or volatile events (Aston-Jones & Cohen, 2005). Phasic NE signaling at a surprising event could lead to higher gain, making the system more responsive to new information that is discordant with the current state, thereby prompting a state transition. In an environment with high volatility or uncertainty, there may be many of these events which are considered surprising or discordant with an active state representation, leading to an elevated tonic state of NE. This tonic state then would decrease the threshold for new salient stimuli to more easily influence the system in a bottom-up manner, driving a positive feedback loop that increases exploratory behavior (Mather, Clewett, Sakaki, & Harley, 2016).

Neural representation of states and transition matrices. In order to effectively partition learning across latent states, the brain must represent these latent states as well as their possible transitions. Investigations of lesions in humans, rodents and monkeys, along with neuroimaging work in humans, has implicated the orbitofrontal cortex (OFC) as a strong candidate for representing and updating latent states. Lesions to the OFC impair reversal learning in a manner that suggests they are relearning after each reversal, rather than recognizing the state transition and retrieving the appropriate alternative policy (Chudasama & Robbins, 2003; Fellows & Farah, 2003; Izquierdo, Suda, & Murray, 2004; Tsuchida, Doll, & Fellows, 2010; Wilson et al., 2014), and careful lesion targeting has localized this function to the OFC's lateral aspects (Rudebeck & Murray, 2011; Rudebeck, Saunders, Lundgren, & Murray, 2017). Human neuroimaging studies have identified latent state representations in OFC (Schuck, Cai, Wilson, & Niv, 2016), and suggest that they transition dynamically at changepoints (Nassar, McGuire, et al., 2019).

OFC is reciprocally connected to the hippocampus, a region that may serve to represent the transition function. Neurons in the rodent hippocampus reflect spatial location according to place fields, but recent work has identified that some such neurons are tuned to represent latent information even within a fixed location (e.g., the number of laps until reward) (Sun, Yang, Martin, & Tonegawa, 2020). Hippocampal activity patterns replay recent spatial transitions in rodents and recent evidence suggests that replay may also occur for more abstract task states in humans (Liu, Dolan, Kurth-Nelson, & Behrens, 2019; Liu, Mattar, Behrens, Daw, & Dolan, 2020; Schuck & Niv, 2019). While the transition function that we defined in our model assigned probabilities to each transition, such probabilities could be estimated by sampling from previous transitions that were experienced. We propose that the replayed patterns of activation in the hippocampus reflect this sampling process, and thus play a role in facilitating appropriate latent state transitions.

As will become apparent in the next section, the OFC and the hippocampus are both regions that play other important roles in the representation of states, including binding information within a context and separating them across contexts.

Segregation of Information Between States

Segregating information between states is key to learning information efficiently, without having to unlearn information of the previous state. Boundaries create well-known effects on memory, disrupting temporal order and distance memory across boundaries (Clewett, Gasser, & Davachi, 2020; DuBrow & Davachi, 2013, 2014, 2016; Horner, Bisby, Wang, Bogus, & Burgess, 2016; Rouhani, Norman, Niv, & Bornstein, 2020). This effect is shown to be the result of disrupting associative binding between items at boundaries (DuBrow & Davachi, 2013), and are marked by pupillometric measures of arousal, much like at change points in adaptive learning paradigms (Clewett et al., 2020). Like the adaptive learning literature, the event segmentation and structure learning literature view the prediction error as key to creating a boundary between states (Gershman et al., 2017; Reynolds, Zacks, & Braver, 2007; Zacks, Speer, Swallow, Braver, & Reynolds, 2007). Even in a paradigm where the probability of transitioning between communities of items in a graph are manipulated to be as equally likely as remaining within the community (Schapiro, Rogers, Cordova, Turk-Browne, & Botvinick, 2013), an agent employing Bayesian inference over efficient representations will still find the transitions between communities to be surprising, creating boundaries between them (Franklin et al., 2020). Latent state models bear some similarities to those of temporal contextual drift and event segmentation (Howard & Kahana, 2002; Reynolds et al., 2007; Zacks et al., 2007). The temporal context drift model posits that a context drifts at a constant rate over time, and events occurring proximally in time are better recalled due to their sharing a similar context (Howard & Kahana, 2002). The model is capable of recognizing episodic events from a previous temporal context; however, it is not capable of rapid updates at changepoints and do not integrate a new occurrence of an event to update parameters of a previous context. Conversely, models of event segmentation are capable of rapid updating when there is a sudden change in context and can account for memory boundary effects (Reynolds et al., 2007; Zacks et al., 2007), but do not have a mechanism to return to a previous state. Latent state models combine the advantages of both temporal context and event segmentation models by bind events by causal states and recognize when these reoccur, capable of rapid adaptation at changepoints (Franklin et al., 2020; Gershman et al., 2017; Gershman & Niv, 2010).

What occurs in the brain to create a “new state”? Place cells in the hippocampus can lend a clue as to the mechanism. These cells encode a specific location in a context (O'Keefe, 1976) and remap between contexts (Muller & Kubie, 1987). What constitutes a change in “context” for the animals empirically depends on a conjunction of factors (O'Keefe & Conway, 1978) and could be described in the same Bayesian inference process described above (Sanders, Wilson, & Gershman, 2020). Interestingly, recent studies find that remapping in place cells is induced by norepinephrine, congruent with the idea that NE signaling during phasic arousal triggers a reset of neural representations at surprising events (Grella, Gomes, Lackie, Renda, & Marrone, 2020; Grella et al., 2019). Thus, it is likely that for non-spatial memories, just as it is for spatial ones, hippocampal cells “remap” when states change (Buzsáki & Moser, 2013). That is not to say that the old representations are lost when a state changes. Indeed, the stored representations of old associations replays in theta cycles and sharp wave ripples, and is recovered when returning to a previous context (Gupta, van der Meer, Touretzky, & Redish, 2010; Liu et al., 2019; Ólafsdóttir, Carpenter, & Barry, 2016). Nevertheless, memories are weaker across boundaries, because they are “chunked” within states to be more efficient (Buzsáki & Moser, 2013).

The interaction between hippocampus and OFC/medial prefrontal cortex (mPFC) is important for binding items within a state, and segregating information across states. During the creation of new associations, the hippocampus completes fast associative inference between connected items while the medial prefrontal cortex extract schematic associations over a longer period of time (Bowman & Zeithamova, 2018; Preston & Eichenbaum, 2013; Schlichting & Preston, 2015; Zeithamova, Dominick, & Preston, 2012; Zeithamova & Preston, 2010). The hippocampus is a good candidate as the initial site of state formation, as its activity spikes right after an event boundary (Baldassano et al., 2017; Ben-Yakov, Eshel, & Dudai, 2013; Ben-Yakov & Henson, 2018; DuBrow & Davachi, 2016), and representations are more similar for within events vs. between events (DuBrow & Davachi, 2016; Schapiro, Turk-Browne, Norman, & Botvinick, 2016). Moreover, activity from hippocampal neurons have been recorded to represent space in segments demarcated by physical landmarks (Gupta, Van Der Meer, Touretzky, & Redish, 2012), and even keep track of higher order structural information in a lap-running task (Sun et al., 2020).

A recent study elegantly showed that hippocampus was necessary for model-based processes only transiently, where a temporary lesion produced a deficit in devaluation only when the procedure occurred a day after initial acquisition, but not a week later (Bradfield et al., 2020). Interestingly, the same study showed that they can similarly disrupt devaluation in the short term by switching the physical context the devaluation takes place, but not in the long term, suggesting that the initial hippocampal dependence was due to reliance on the spatial context. The OFC, which is critical for devaluation (Rudebeck and Murray, 2011; Rudebeck et al., 2017; Reber et al., 2017), is a good candidate for the computation of this type of higher order association at longer time horizons, free of the physical context – which would be interesting to test formally.

The available evidence suggests that the mPFC and OFC is key to maintaining the relevant state, in order to disambiguate similar associations that exist across different states (Wilson et al., 2014). In studies with functional neuroimaging, the mPFC has been shown to increase in functional connectivity with the hippocampus within event boundaries (Dubrow & Davachi, 2016), and exhibit consistent activity across stories with the same script (e.g., eating at a restaurant) but different characters and details, even across different modalities (Baldassano, Hasson, & Norman, 2018). Consistent with this, lesions of the OFC/mPFC region disrupt script and schema knowledge in humans (Ghosh, Moscovitch, Colella, & Gilboa, 2014; Spalding, Jones, Duff, Tranel, & Warren, 2015; Wood, Tierney, Bidwell, & Grafman, 2005). In rats, goal-related mPFC activity were closely entrained to hippocampal theta rhythm (Hyman, Zilli, Paley, & Hasselmo, 2005), and inactivation of the mPFC disrupts it from providing contextual information to hippocampus and increases interference between memories from different contexts (Guise & Shapiro, 2017).

The mPFC and OFC then passes the state and action mappings through direct connections to ventral striatum, the dopaminergic projections from which perform reinforcement learning over actions (Frank et al., 2004; Kravitz, Tye, & Kreitzer, 2012). Through the corticostriatal circuit between the OFC/mPFC, striatum, and thalamus, the associations between state and actions are reinforced or weakened (Haber, 2011). At the same time, other connections from the hippocampus (e.g., through the subiculum and entorhinal cortex) to the striatum also could convey state information for reinforcement learning.

Figure 3. hippocampal remapping vs. entorhinal reorientation. When you go to your friend’s place after the pandemic is over, you are surprised to find that your friend has redecorated her living room in a fit of boredom during the quarantine. We propose that this surprise is communicated to multiple brain regions through a transient norepinephrine signal, and has different effects on neural representations in different brain regions. The place cells that fired in specific areas (firing fields of 3 example place cells in red, green and blue) in her living room will remap after the surprise signal to different places, which enables fast unlearning of previous behaviors or associations (e.g., learned behaviors that are stored in associations to a population of neurons with similar spatial tuning will not be elicited when that population of neurons is re-distributed to tile non-overlapping locations). Entorhinal grid cells (a subset of the firing field of a single grid cell in yellow), on the other hand, preserve the structure of the relationships in the room: the firing fields of the cell over the TV set and the couch keep their relative relationship across the rotation, so the same highly valued action of sitting on the couch that faces the TV still holds true despite the location change of these items.

Transfer of Learning Between Latent States

While segregation of information is important to prevent sharing of irrelevant information, it is desirable to share certain types of information across different states. Such information is often the *structure* of the task, defined as low dimensional relationships among entities that can be abstracted out of individual experiences and generalized to new situations. A good example is that when you go to a new grocery store in a different city, or even country, and you are trying to find the dairy aisle, you may not need to learn the layout of the grocery store from scratch. Instead, you can draw upon your abstracted knowledge of the many grocery

stores you have visited over the course of your life, and realize that the dairy aisle is probably opposite to the produce, and next to the meat section. Generalization is a hard problem for artificial agents, with even sophisticated meta-learning agents achieving generalization only to very similar tasks (Wang et al., 2016).

Some insight into how structure is shared across contexts might be gleaned by considering how structure is shared within contexts. One computational strategy for sharing information across similar situations is through the use of basis functions. A basis set provides a way of re-representing a continuous state space through the linear combination of individual basis functions (e.g., neurons) that might be “tuned” to different features of the environment. Learning in this framework involves learning the weights that map each basis function onto a quantity of interest, for example, value. A linear basis set for functional approximation of expected value function of a state would be equal to the sum of weights multiplied by the feature vector (Sutton & Barto, 1998). The specific choice of basis functions in our basis set will control how information is shared across different situations – if two situations are represented by activation of a single basis function (e.g., neuron) then weights associated to them will apply to both of these situations. In the same spirit, if our basis set represented all situations similarly in two different contexts – any learning done in one context would be transferred to the other. In principle, more complex transfer of knowledge across contexts could be possible if basis functions had more complex relationships across contexts, for example, if all basis functions were rotated coherently in physical space in one context relative to the other.

One candidate neural mechanism for achieving this kind of transfer of structural information across contexts is the entorhinal grid cell system (Hafting, Fyhn, Molden, Moser, & Moser, 2005). Grid cell firing fields form tessellating hexagonal tiling over an animal’s current environment, and, unlike hippocampal place cells, which remap between environments, grid cells’ firing fields realign in concert in a new environment while maintaining their relative positions to each other (Fyhn, Hafting, Treves, Moser, & Moser, 2007). In this way, grid cells act as a tiled feature vector (given a uniform space where all directions are possible – in spaces where movements are restricted in certain directions the regularity breaks down (Stachenfeld, Botvinick, & Gershman, 2017). This consistency could be key to maintaining structural information across contexts. Beyond the entorhinal cortex itself, however, other brain regions, including the mPFC, have shown grid-like coding to both spatial and abstract cognitive variables (Constantinescu, O’Reilly, & Behrens, 2016; Doeller, Barry, & Burgess, 2010; Jacobs et al., 2013; Park, Miller, & Boorman, 2020). Thus, the grid code may be a general principle of a type of information coding, present throughout the brain, important not just in navigation of physical space, but also in abstract concepts and memories (Behrens et al., 2018; Bellmund, Gärdenfors, Moser, & Doeller, 2018; Whittington et al., 2020).

There are many theories about how the grid code could enable generalization and transfer of information both within and across contexts. The grid code has been proposed to represent a 2D code of cognitive variables that combine to scale up to higher dimensions, or enable novel inferences of relationships by performing vector navigation (Bush, Barry, Manson, & Burgess, 2015; Klukas, Lewis, & Fiete, 2020). It has been observed that the grid cell could be expressed as a principal component of place cell firing fields (Behrens et al., 2018; Dordek, Soudry, Meir, & Derdikman, 2016; Stachenfeld et al., 2017). One account we will highlight here is that the grid code might reflect an eigenvector of the successor representation (SR;

Stachenfeld et al., 2017), because it comes with some interesting implications and constraints. A successor representation is a cached representation of the transition structure of a space, learnt while an agent traversed the space under a certain policy (Dayan, 1993). Taking the grocery store example, let's say that the agent learns that making a right at the registers and then turning left leads it to the ice cream freezer. SR enables a type of planning without the more effortful and thorough kind usually required under model-based strategies such as dynamic programming (Sutton & Barto, 1998). If the store decided to move the ice cream to a different freezer somewhere along the same route the agent has traveled before, the agent will be able to immediately reroute to the new location with the SR without further training (Momennejad et al., 2017). However, if the ice cream is moved to a part of the store the agent had never explored under the previous policy, then the SR would not help the agent find it. This account, then, implies that inference of novel routes would not be able to be performed by grid cells, and that planning itself must operate elsewhere (eg. simulation/replay in the hippocampus)(Mattar & Daw, 2018), which runs counter to some other theories of grid cell function (Bush et al., 2015). However, taking the principal components of the successor representation *would* enable fast credit assignment, such that, when applied to a new environment with a similar structure (e.g., a different grocery store), information about the relationships of different items within that layout is preserved and highly valued action contingencies are able to be immediately applied without further training – a result that does not immediately fall out of the predictions of other accounts of grid cells.

Accordingly, we propose that, when a highly surprising observation occurs, the norepinephrine signal would prompt remapping in hippocampal to promote fast disengagement from previously learned action associations, but lead to the realignment of grid codes to promote fast transfer of action values that are relevant to the new environment. Walking through an example, let's say that when you go to your friend's house after the pandemic is over, you are surprised to see she has redecorated in a fit of quarantine-induced boredom (Figure 3). The bookshelf has been moved to the opposite wall, and the couch and TV stand has been moved to a different part of the room as well. In this situation, your place cells for your friend's living room would likely remap. This remapping is useful for preventing interference; if all place cells located near the couch in the first context were co-localized in the second one, learned associations might lead you to embarrass yourself by attempting to sit where there is no couch. Remapping, from this view, serves to break the associations that were attributed to specific locations and to instead, distribute them across the environment (or remove them from the new context completely)(Alme et al., 2014). On the hand, your grid cells – more precisely, the modules of the grid cells corresponding to the spatial scale of the room (Stensola et al., 2012) -- will realign in concert, preserving the relative relationships between each other (Fyhn et al., 2007). This coordinated rotation preserves the relationship between the items that can be useful for transferring valued actions. For instance, the relationship between the couch and the TV is a highly relevant one for your goal of watching a movie with your friend. The reorientation of the grid cells preserves the action values of the TV-couch relationship and transfers it to the new room arrangement, so that no relearning is required. This could allow a principled way to transfer value gradients from one environment to another.

The entorhinal system, which also includes cells that mark borders and head direction as well as grid cells, provides the input to the hippocampus in forming place cells (though it is

important to note the reciprocal interdependence between the two systems) (Bonnievie et al., 2013; Moser, Rowland, & Moser, 2015). Because of the influence of grid cells on place cells, several accounts have suggested that remapping of place cells are, in fact, not completely random, but are governed by the reorientation of several input grid modules (Moser et al., 2015) or to the phase of the grid code (Whittington et al., 2020). Accordingly, the entorhinal grid system could provide the stable framework that permits the arrangement of sensory information during reconsolidation and planning. Replay of hippocampal place cells occur in concert with entorhinal grid cells during rest (Ólafsdóttir et al., 2016) and hippocampal replay during quiet awake mode is found to be significantly more fragmented after inhibition of inputs from the medial entorhinal cortex (Yamamoto & Tonegawa, 2017). The constancy of the dimensional structure in the entorhinal cortex is maintained through distortions, even while hippocampal and sensory inputs are changed (Yoon et al., 2013). The specific purposes of the grid code in mPFC and other parts of the cortex has not so far been studied, but it is possible that inputs from the entorhinal cortex enforces structure on state representation and helps to create a state-to-state transition map, which can be used for goal-relevant action planning.

Summary

When changes happen, how we react to it depends on our inference of the underlying context. In this review, we have summarized insights gained from adaptive learning and structural learning, to describe how the brain can recognize a new state, segregate information between states, and transfer relevant structural information across states. We hypothesize that mismatch between the evidence and top-down expectations from an internal model produce a surprise signal, which prompts the phasic release of norepinephrine. This promotes a state transition, which is initiated by remapping in the hippocampus and maintained to goal-relevance by the medial prefrontal cortex. The entorhinal system extracts low-dimensional, structural information from the state, and transfers it to new contexts to enable fast generalization. The goal-relevant state in the orbitofrontal cortex, as well as state information from other parts of the cortex and the hippocampal complex, is passed to the striatum, which conducts reinforcement learning that strengthens or weakens the actions based on consequences within the corticostriatal loop.

Future directions

We have laid out a schematic for how adaptive learning and latent state learning work together to select relevant actions in a changing environment. However, many of the mechanisms within it have yet to be hashed out, which can be summarized as two major questions. First, though it is clear that efficient learning strategies in complex environments require crediting latent states that are inferred, rather than directly visible, the exact mechanisms through which the brain makes such inferences remain somewhat unclear. Neuroimaging and lesion evidence point to both the medial prefrontal cortex and hippocampus as important for inference and generalization (Bowman & Zeithamova, 2018; Pajkert et al., 2017; Spalding et al., 2018; Zeithamova et al., 2012), but the manner in which they do so is debated. Particularly, what is the code used to make the connection between two things that are learnt separately? Is it through pattern completion forming an integrated memory trace,

inference from abstract categories, a sharing of common neurons in a distributed code, or spreading activation through recurrent networks (Bowman & Zeithamova, 2018; Greene, Gross, Elsinger, & Rao, 2006; Kumaran & McClelland, 2012; Schapiro, Turk-Browne, Botvinick, & Norman, 2017; Shohamy & Wagner, 2008; Zeithamova et al., 2012)? Furthermore, what is the role of the grid code in this process? Is it one that would allow for inference of novel associations on its own akin to vector navigation (Bush et al., 2015), or does it instead provide a low-dimensional summary of already learned associations (Stachenfeld et al., 2017) for faster credit assignment in new situations, as we propose here? Is it used in egocentric action planning, or allocentric policy planning, or does that depend on which area of the brain (e.g., mPFC vs. entorhinal vs. posterior cingulate cortex) is using the code? Does it represent a positional code that enables fast generalization to any item or concept swapped within those positions (Whittington et al., 2020), or does it instead encode relationships between concepts?

The second major question yet to be answered relates to how dynamic signals control which representations are active, in order to give proper credit assignment. We have seen that phasic NE can induce place cell remapping in the hippocampus (Grella et al., 2020; Grella et al., 2019), but is it sufficient to produce reorientation in grid cells? Does it prompt the loading task states or transition to specific task dimensions? When states are interleaved (for example the positive and negative outcomes that are associated with separate learning rates in Pulcu & Browning, 2017), how does the brain switch back and forth rapidly and assign appropriate credit to each? Finally, when there is ambiguity during learning and credit is wrongly assigned – can the brain go back and resolve the error later on? Models of reinforcement learning suggests that simulation of experienced events would expedite learning (Sutton & Barto, 1998), which is supported by the evidence of replay events in the hippocampus. Incredibly, this replay can even arrange disjointed events in their correct order (Liu et al., 2019). Yet, it is still not very well understood how exactly it is we learn from these replay events. Though new work has shown behaviorally that a model-based system can guide credit assignment and correct errors, the neural mechanisms of this process is not known (Moran, Keramati, Dayan, & Dolan, 2019). Do NE and dopaminergic signals play a role in learning from replay as well?

The question of how the brain identifies states, makes inferences of latent information, and learns from relevant sources is at the forefront of cognitive neuroscience. Artificial intelligence train from hundreds of thousands of exemplars, and can only generalize their learning in the best cases to highly similar tasks. Humans, on the other hand, sometimes need only a single example in order to make the correct generalization. Learning the mechanism underlying these functions are not only key to understanding the most advanced and unique aspects of human cognition, but also to produce artificial intelligence that can quickly learn and generalize abstract principles.

References

Adams, R. P., & MacKay, D. J. (2007). *Bayesian online changepoint detection*. Cambridge Cambridge University.

- Alme, C. B., Miao, C., Jezek, K., Treves, A., Moser, E. I., & Moser, M.-B. (2014). Place cells in the hippocampus: eleven maps for eleven rooms. *Proceedings of the National Academy of Sciences*, 111(52), 18428-18435.
- Aston-Jones, G., & Bloom, F. (1981). Activity of norepinephrine-containing locus coeruleus neurons in behaving rats anticipates fluctuations in the sleep-waking cycle. *Journal of Neuroscience*, 1(8), 876-886.
- Aston-Jones, G., & Cohen, J. D. (2005). An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance. *Annu. Rev. Neurosci.*, 28, 403-450.
- Baldassano, C., Chen, J., Zadbood, A., Pillow, J. W., Hasson, U., & Norman, K. A. (2017). Discovering event structure in continuous narrative perception and memory. *Neuron*, 95(3), 709-721. e705.
- Baldassano, C., Hasson, U., & Norman, K. A. (2018). Representation of real-world event schemas during narrative perception. *Journal of Neuroscience*, 38(45), 9689-9699.
- Behrens, T. E., Muller, T. H., Whittington, J. C., Mark, S., Baram, A. B., Stachenfeld, K. L., & Kurth-Nelson, Z. (2018). What is a cognitive map? Organizing knowledge for flexible behavior. *Neuron*, 100(2), 490-509.
- Behrens, T. E., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, 10(9), 1214-1221.
- Bellmund, J. L., Gärdenfors, P., Moser, E. I., & Doeller, C. F. (2018). Navigating cognition: Spatial codes for human thinking. *Science*, 362(6415), eaat6766.
- Ben-Yakov, A., Eshel, N., & Dudai, Y. (2013). Hippocampal immediate poststimulus activity in the encoding of consecutive naturalistic episodes. *Journal of Experimental Psychology: General*, 142(4), 1255.
- Ben-Yakov, A., & Henson, R. N. (2018). The Hippocampal Film Editor: Sensitivity and Specificity to Event Boundaries in Continuous Experience. *The Journal of Neuroscience*, 38(47), 10057-10068. doi:10.1523/jneurosci.0524-18.2018
- Bonnevie, T., Dunn, B., Fyhn, M., Hafting, T., Derdikman, D., Kubie, J. L., . . . Moser, M.-B. (2013). Grid cells require excitatory drive from the hippocampus. *Nature Neuroscience*, 16(3), 309-317. doi:10.1038/nn.3311
- Bouret, S., & Sara, S. J. (2005). Network reset: a simplified overarching theory of locus coeruleus noradrenaline function. *Trends in neurosciences*, 28(11), 574-582.
- Bowman, C. R., & Zeithamova, D. (2018). Abstract memory representations in the ventromedial prefrontal cortex and hippocampus support concept generalization. *Journal of Neuroscience*, 38(10), 2605-2614.
- Browning, M., Behrens, T. E., Jocham, G., O'Reilly, J. X., & Bishop, S. J. (2015). Anxious individuals have difficulty learning the causal statistics of aversive environments. *Nature Neuroscience*, 18(4), 590-596.
- Bush, D., Barry, C., Manson, D., & Burgess, N. (2015). Using grid cells for navigation. *Neuron*, 87(3), 507-520.
- Buzsáki, G., & Moser, E. I. (2013). Memory, navigation and theta rhythm in the hippocampal-entorhinal system. *Nature Neuroscience*, 16(2), 130-138.
- Chudasama, Y., & Robbins, T. W. (2003). Dissociable contributions of the orbitofrontal and infralimbic cortex to pavlovian autoshaping and discrimination reversal learning: further

- evidence for the functional heterogeneity of the rodent frontal cortex. *Journal of Neuroscience*, 23(25), 8771-8780.
- Clewett, D., Gasser, C., & Davachi, L. (2020). Pupil-linked arousal signals track the temporal organization of events in memory. *Nature communications*, 11(1), 1-14.
- Collins, A. G., & Frank, M. J. (2013). Cognitive control over learning: Creating, clustering, and generalizing task-set structure. *Psychological Review*, 120(1), 190.
- Collins, A. G., & Koechlin, E. (2012). Reasoning, learning, and creativity: frontal lobe function and human decision-making. *PLoS biology*, 10(3), e1001293.
- Constantinescu, A. O., O'Reilly, J. X., & Behrens, T. E. (2016). Organizing conceptual knowledge in humans with a gridlike code. *Science*, 352(6292), 1464-1468.
- d'Acremont, M., & Bossaerts, P. (2016). Neural mechanisms behind identification of leptokurtic noise and adaptive behavioral response. *Cerebral Cortex*, 26(4), 1818-1830.
- Dayan, P. (1993). Improving Generalization for Temporal Difference Learning: The Successor Representation. *Neural computation*, 5(4), 613-624. doi:10.1162/neco.1993.5.4.613
- Doeller, C. F., Barry, C., & Burgess, N. (2010). Evidence for grid cells in a human memory network. *Nature*, 463(7281), 657.
- Donoso, M., Collins, A. G., & Koechlin, E. (2014). Foundations of human reasoning in the prefrontal cortex. *Science*, 344(6191), 1481-1486.
- Dordek, Y., Soudry, D., Meir, R., & Derdikman, D. (2016). Extracting grid cell characteristics from place cell inputs using non-negative principal component analysis. *eLife*, 5, e10094.
- DuBrow, S., & Davachi, L. (2013). The influence of context boundaries on memory for the sequential order of events. *Journal of Experimental Psychology: General*, 142(4), 1277.
- DuBrow, S., & Davachi, L. (2014). Temporal memory is shaped by encoding stability and intervening item reactivation. *Journal of Neuroscience*, 34(42), 13998-14005.
- DuBrow, S., & Davachi, L. (2016). Temporal binding within and across events. *Neurobiology of Learning and Memory*, 134, 107-114.
- Ebitz, R. B., Albarran, E., & Moore, T. (2018). Exploration disrupts choice-predictive signals and alters dynamics in prefrontal cortex. *Neuron*, 97(2), 450-461. e459.
- Fellows, L. K., & Farah, M. J. (2003). Ventromedial frontal cortex mediates affective shifting in humans: evidence from a reversal learning paradigm. *Brain : a journal of neurology*, 126(Pt 8), 1830-1837. doi:10.1093/brain/awg180
- Filipowicz, A. L., Glaze, C. M., Kable, J. W., & Gold, J. I. (2020). Pupil diameter encodes the idiosyncratic, cognitive complexity of belief updating. *eLife*, 9, e57872.
- Fischer, A. G., & Ullsperger, M. (2013). Real and fictive outcomes are processed differently but converge on a common adaptive mechanism. *Neuron*, 79(6), 1243-1255.
- Foote, S., Aston-Jones, G., & Bloom, F. (1980). Impulse activity of locus coeruleus neurons in awake rats and monkeys is a function of sensory stimulation and arousal. *Proceedings of the National Academy of Sciences*, 77(5), 3033-3037.
- Frank, M. J., Seeberger, L. C., & O'Reilly, R. C. (2004). By carrot or by stick: cognitive reinforcement learning in parkinsonism. *Science*, 306(5703), 1940-1943.
- Franklin, N. T., Norman, K. A., Ranganath, C., Zacks, J. M., & Gershman, S. J. (2020). Structured Event Memory: A neuro-symbolic model of event cognition. *Psychological Review*, 127(3), 327.

- Friston, K. (2012). The history of the future of the Bayesian brain. *NeuroImage*, 62(2), 1230-1233. doi:10.1016/j.neuroimage.2011.10.004
- Fyhn, M., Hafting, T., Treves, A., Moser, M.-B., & Moser, E. I. (2007). Hippocampal remapping and grid realignment in entorhinal cortex. *Nature*, 446(7132), 190-194.
- Gershman, S. J., Blei, D. M., & Niv, Y. (2010). Context, learning, and extinction. *Psychological Review*, 117(1), 197.
- Gershman, S. J., Jones, C. E., Norman, K. A., Monfils, M.-H., & Niv, Y. (2013). Gradual extinction prevents the return of fear: implications for the discovery of state. *Frontiers in behavioral neuroscience*, 7, 164.
- Gershman, S. J., Monfils, M.-H., Norman, K. A., & Niv, Y. (2017). The computational nature of memory modification. *eLife*, 6, e23763.
- Gershman, S. J., & Niv, Y. (2010). Learning latent structure: carving nature at its joints. *Current Opinion in Neurobiology*, 20(2), 251-256.
- Gershman, S. J., Radulescu, A., Norman, K. A., & Niv, Y. (2014). Statistical computations underlying the dynamics of memory updating. *PLoS Comput Biol*, 10(11), e1003939.
- Ghosh, V. E., Moscovitch, M., Colella, B. M., & Gilboa, A. (2014). Schema representation in patients with ventromedial PFC lesions. *The Journal of Neuroscience*, 34(36), 12057-12070.
- Greene, A. J., Gross, W. L., Elsinger, C. L., & Rao, S. M. (2006). An fMRI analysis of the human hippocampus: inference, context, and task awareness. *Journal of cognitive neuroscience*, 18(7), 1156-1173.
- Grella, S. L., Gomes, S. M., Lackie, R. E., Renda, B., & Marrone, D. F. (2020). Norepinephrine as a Memory Reset Signal: Switching the System from Retrieval to Encoding During a Spatial Memory Task can be Both Adaptive and Maladaptive. *bioRxiv*.
- Grella, S. L., Neil, J. M., Edison, H. T., Strong, V. D., Odintsova, I. V., Walling, S. G., . . . Harley, C. W. (2019). Locus coeruleus phasic, but not tonic, activation initiates global remapping in a familiar environment. *Journal of Neuroscience*, 39(3), 445-455.
- Guise, K. G., & Shapiro, M. L. (2017). Medial prefrontal cortex reduces memory interference by modifying hippocampal encoding. *Neuron*, 94(1), 183-192. e188.
- Gupta, A. S., van der Meer, M. A., Touretzky, D. S., & Redish, A. D. (2010). Hippocampal replay is not a simple function of experience. *Neuron*, 65(5), 695-705.
- Gupta, A. S., Van Der Meer, M. A., Touretzky, D. S., & Redish, A. D. (2012). Segmentation of spatial experience by hippocampal theta sequences. *Nature Neuroscience*, 15(7), 1032-1039.
- Haber, S. N. (2011). Neural circuits of reward and decision making: Integrative networks across corticobasal ganglia loops. *Neural basis of motivational and cognitive control*, 21-35.
- Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436(7052), 801-806.
- Heald, J. B., Lengyel, M., & Wolpert, D. M. (2020). Contextual inference underlies the learning of sensorimotor repertoires. *bioRxiv*, 2020.2011.2023.394320. doi:10.1101/2020.11.23.394320
- Horner, A. J., Bisby, J. A., Wang, A., Bogus, K., & Burgess, N. (2016). The role of spatial boundaries in shaping long-term event representations. *Cognition*, 154, 151-164.

- Howard, M. W., & Kahana, M. J. (2002). A distributed representation of temporal context. *Journal of Mathematical Psychology*, 46(3), 269-299.
- Hyman, J. M., Zilli, E. A., Paley, A. M., & Hasselmo, M. E. (2005). Medial prefrontal cortex cells show dynamic modulation with the hippocampal theta rhythm dependent on behavior. *Hippocampus*, 15(6), 739-749.
- Izquierdo, A., Suda, R. K., & Murray, E. A. (2004). Bilateral orbital prefrontal cortex lesions in rhesus monkeys disrupt choices guided by both reward value and reward contingency. *Journal of Neuroscience*, 24(34), 7540-7548.
- Jacobs, J., Weidemann, C. T., Miller, J. F., Solway, A., Burke, J. F., Wei, X.-X., . . . Fried, I. (2013). Direct recordings of grid-like neuronal activity in human spatial navigation. *Nature Neuroscience*, 16(9), 1188-1190.
- Jepma, M., Murphy, P. R., Nassar, M. R., Rangel-Gomez, M., Meeter, M., & Nieuwenhuis, S. (2016). Catecholaminergic regulation of learning rate in a dynamic environment. *PLoS computational biology*, 12(10), e1005171.
- Joshi, S., Li, Y., Kalwani, R. M., & Gold, J. I. (2016). Relationships between pupil diameter and neuronal activity in the locus coeruleus, colliculi, and cingulate cortex. *Neuron*, 89(1), 221-234.
- Kao, C.-H., Khambhati, A. N., Bassett, D. S., Nassar, M. R., McGuire, J. T., Gold, J. I., & Kable, J. W. (2020). Functional brain network reconfiguration during learning in a dynamic environment. *Nature communications*, 11(1), 1-13.
- Klukas, M., Lewis, M., & Fiete, I. (2020). Efficient and flexible representation of higher-dimensional cognitive variables with grid cells. *PLoS computational biology*, 16(4), e1007796.
- Knill, D. C., & Richards, W. (1996). *Perception as Bayesian inference*: Cambridge University Press.
- Kravitz, A. V., Tye, L. D., & Kreitzer, A. C. (2012). Distinct roles for direct and indirect pathway striatal neurons in reinforcement. *Nature Neuroscience*, 15(6), 816-818.
- Krishnamurthy, K., Nassar, M. R., Sarode, S., & Gold, J. I. (2017). Arousal-related adjustments of perceptual biases optimize perception in dynamic environments. *Nature human behaviour*, 1(6), 1-11.
- Kumaran, D., & McClelland, J. L. (2012). Generalization through the recurrent interaction of episodic memories: a model of the hippocampal system. *Psychological Review*, 119(3), 573.
- Liu, Y., Dolan, R. J., Kurth-Nelson, Z., & Behrens, T. E. (2019). Human replay spontaneously reorganizes experience. *Cell*, 178(3), 640-652. e614.
- Liu, Y., Mattar, M. G., Behrens, T. E. J., Daw, N. D., & Dolan, R. J. (2020). Experience replay supports non-local learning. *bioRxiv*, 2020.2010.2020.343061. doi:10.1101/2020.10.20.343061
- Lloyd, K., & Leslie, D. S. (2013). Context-dependent decision-making: a simple Bayesian model. *Journal of The Royal Society Interface*, 10(82), 20130069.
- Maes, E. J., Sharpe, M. J., Usypchuk, A. A., Lozzi, M., Chang, C. Y., Gardner, M. P., . . . Iordanova, M. D. (2020). Causal evidence supporting the proposal that dopamine transients function as temporal difference prediction errors. *Nature Neuroscience*, 23(2), 176-178.

- Mather, M., Clewett, D., Sakaki, M., & Harley, C. W. (2016). Norepinephrine ignites local hotspots of neuronal excitation: How arousal amplifies selectivity in perception and memory. *Behavioral and Brain Sciences*, 39.
- Mattar, M. G., & Daw, N. D. (2018). Prioritized memory access explains planning and hippocampal replay. *Nature Neuroscience*, 21(11), 1609-1617.
- McGuire, J. T., Nassar, M. R., Gold, J. I., & Kable, J. W. (2014). Functionally dissociable influences on learning rate in a dynamic environment. *Neuron*, 84(4), 870-881.
- Meyniel, F., & Dehaene, S. (2017). Brain networks for confidence weighting and hierarchical inference during probabilistic learning. *Proceedings of the National Academy of Sciences*, 114(19), E3859-E3868.
- Momennejad, I., Russek, E. M., Cheong, J. H., Botvinick, M. M., Daw, N. D., & Gershman, S. J. (2017). The successor representation in human reinforcement learning. *Nature human behaviour*, 1(9), 680-692.
- Moran, R., Keramati, M., Dayan, P., & Dolan, R. J. (2019). Retrospective model-based inference guides model-free credit assignment. *Nature communications*, 10(1), 750. doi:10.1038/s41467-019-08662-8
- Moser, M.-B., Rowland, D. C., & Moser, E. I. (2015). Place cells, grid cells, and memory. *Cold Spring Harbor perspectives in biology*, 7(2), a021808.
- Muller, R. U., & Kubie, J. L. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *Journal of Neuroscience*, 7(7), 1951-1968.
- Nassar, M. R., Bruckner, R., & Frank, M. J. (2019). Statistical context dictates the relationship between feedback-related EEG signals and learning. *eLife*, 8, e46975.
- Nassar, M. R., McGuire, J. T., Ritz, H., & Kable, J. W. (2019). Dissociable forms of uncertainty-driven representational change across the human brain. *Journal of Neuroscience*, 39(9), 1688-1698.
- Nassar, M. R., Rumsey, K. M., Wilson, R. C., Parikh, K., Heasley, B., & Gold, J. I. (2012). Rational regulation of learning dynamics by pupil-linked arousal systems. *Nature Neuroscience*, 15(7), 1040.
- Nour, M. M., Dahoun, T., Schwartenbeck, P., Adams, R. A., FitzGerald, T. H., Coello, C., . . . Howes, O. D. (2018). Dopaminergic basis for signaling belief updates, but not surprise, and the link to paranoia. *Proceedings of the National Academy of Sciences*, 115(43), E10167-E10176.
- O'Keefe, J. (1976). Place units in the hippocampus of the freely moving rat. *Experimental neurology*, 51(1), 78-109.
- O'Keefe, J., & Conway, D. (1978). Hippocampal place units in the freely moving rat: why they fire where they fire. *Experimental Brain Research*, 31(4), 573-590.
- O'Reilly, J. X., Schüffelgen, U., Cuell, S. F., Behrens, T. E., Mars, R. B., & Rushworth, M. F. (2013). Dissociable effects of surprise and model update in parietal and anterior cingulate cortex. *Proceedings of the National Academy of Sciences*, 110(38), E3660-E3669.
- Ólafsdóttir, H. F., Carpenter, F., & Barry, C. (2016). Coordinated grid and place cell replay during rest. *Nature Neuroscience*, 19(6), 792-794.
- Pajkert, A., Finke, C., Shing, Y. L., Hoffmann, M., Sommer, W., Heekeren, H. R., & Ploner, C. J. (2017). Memory integration in humans with hippocampal lesions. *Hippocampus*, 27(12), 1230-1238. doi:10.1002/hipo.22766

- Park, S. A., Miller, D. S., & Boorman, E. D. (2020). Novel Inferences in a Multidimensional Social Network Use a Grid-like Code. *bioRxiv*.
- Preston, A. R., & Eichenbaum, H. (2013). Interplay of hippocampus and prefrontal cortex in memory. *Current Biology*, 23(17), R764-R773.
- Pulcu, E., & Browning, M. (2017). Affective bias as a rational response to the statistics of rewards and punishments. *eLife*, 6, e27879.
- Razmi, N., & Nassar, M. R. (2020). Adaptive learning through temporal dynamics of state representation. *bioRxiv*.
- Reimer, J., McGinley, M. J., Liu, Y., Rodenkirch, C., Wang, Q., McCormick, D. A., & Tolias, A. S. (2016). Pupil fluctuations track rapid changes in adrenergic and cholinergic activity in cortex. *Nature communications*, 7(1), 1-7.
- Reynolds, J. R., Zacks, J. M., & Braver, T. S. (2007). A computational model of event segmentation from perceptual prediction. *Cognitive science*, 31(4), 613-643.
- Rouhani, N., Norman, K. A., Niv, Y., & Bornstein, A. M. (2020). Reward prediction errors create event boundaries in memory. *Cognition*, 203, 104269.
- Rudebeck, P. H., & Murray, E. A. (2011). Dissociable effects of subtotal lesions within the macaque orbital prefrontal cortex on reward-guided behavior. *Journal of Neuroscience*, 31(29), 10569-10578.
- Rudebeck, P. H., Saunders, R. C., Lundgren, D. A., & Murray, E. A. (2017). Specialized representations of value in the orbital and ventrolateral prefrontal cortex: desirability versus availability of outcomes. *Neuron*, 95(5), 1208-1220. e1205.
- Sanders, H., Wilson, M. A., & Gershman, S. J. (2020). Hippocampal remapping as hidden state inference. *eLife*, 9, e51140.
- Schapiro, A. C., Rogers, T. T., Cordova, N. I., Turk-Browne, N. B., & Botvinick, M. M. (2013). Neural representations of events arise from temporal community structure. *Nature Neuroscience*, 16(4), 486-492.
- Schapiro, A. C., Turk-Browne, N. B., Botvinick, M. M., & Norman, K. A. (2017). Complementary learning systems within the hippocampus: a neural network modelling approach to reconciling episodic memory with statistical learning. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 372(1711), 20160049. doi:doi:10.1098/rstb.2016.0049
- Schapiro, A. C., Turk-Browne, N. B., Norman, K. A., & Botvinick, M. M. (2016). Statistical learning of temporal community structure in the hippocampus. *Hippocampus*, 26(1), 3-8.
- Schlichting, M. L., & Preston, A. R. (2015). Memory integration: neural mechanisms and implications for behavior. *Current Opinion in Behavioral Sciences*, 1, 1-8.
- Schuck, N. W., Cai, M. B., Wilson, R. C., & Niv, Y. (2016). Human orbitofrontal cortex represents a cognitive map of state space. *Neuron*, 91(6), 1402-1412.
- Schuck, N. W., & Niv, Y. (2019). Sequential replay of nonspatial task states in the human hippocampus. *Science*, 364(6447), eaaw5181.
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A Neural Substrate of Prediction and Reward. *Science*, 275(5306), 1593-1599. doi:10.1126/science.275.5306.1593
- Shohamy, D., & Wagner, A. D. (2008). Integrating memories in the human brain: hippocampal-midbrain encoding of overlapping events. *Neuron*, 60(2), 378-389.

- Spalding, K. N., Jones, S. H., Duff, M. C., Tranel, D., & Warren, D. E. (2015). Investigating the Neural Correlates of Schemas: Ventromedial Prefrontal Cortex Is Necessary for Normal Schematic Influence on Memory. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 35(47), 15746-15751. doi:10.1523/JNEUROSCI.2767-15.2015
- Spalding, K. N., Schlichting, M. L., Zeithamova, D., Preston, A. R., Tranel, D., Duff, M. C., & Warren, D. E. (2018). Ventromedial prefrontal cortex is necessary for normal associative inference and memory integration. *Journal of Neuroscience*, 38(15), 3767-3775.
- Stachenfeld, K. L., Botvinick, M. M., & Gershman, S. J. (2017). The hippocampus as a predictive map. *Nature Neuroscience*, 20(11), 1643-1653. doi:10.1038/nn.4650
- Starkweather, C. K., Babayan, B. M., Uchida, N., & Gershman, S. J. (2017). Dopamine reward prediction errors reflect hidden-state inference across time. *Nature Neuroscience*, 20(4), 581-589.
- Stensola, H., Stensola, T., Solstad, T., Frøland, K., Moser, M.-B., & Moser, E. I. (2012). The entorhinal grid map is discretized. *Nature*, 492(7427), 72-78.
- Sun, C., Yang, W., Martin, J., & Tonegawa, S. (2020). Hippocampal neurons represent events as transferable units of experience. *Nature Neuroscience*, 23(5), 651-663.
- Sutton, R. S., & Barto, A. G. (1998). *Introduction to reinforcement learning* (Vol. 135): MIT press Cambridge.
- Tsuchida, A., Doll, B. B., & Fellows, L. K. (2010). Beyond reversal: A critical role for human orbitofrontal cortex in flexible learning from probabilistic feedback. *The Journal of Neuroscience*, 30(50), 16868-16875. doi:10.1523/JNEUROSCI.1958-10.2010
- Vazey, E. M., Moorman, D. E., & Aston-Jones, G. (2018). Phasic locus coeruleus activity regulates cortical encoding of salience information. *Proceedings of the National Academy of Sciences*, 115(40), E9439-E9448.
- Waelti, P., Dickinson, A., & Schultz, W. (2001). Dopamine responses comply with basic assumptions of formal learning theory. *Nature*, 412(6842), 43-48.
- Wang, J. X., Kurth-Nelson, Z., Tirumala, D., Soyer, H., Leibo, J. Z., Munos, R., . . . Botvinick, M. (2016). Learning to reinforcement learn. *arXiv preprint arXiv:1611.05763*.
- Whittington, J. C., Muller, T. H., Mark, S., Chen, G., Barry, C., Burgess, N., & Behrens, T. E. (2020). The Tolman-Eichenbaum machine: Unifying space and relational memory through generalization in the hippocampal formation. *Cell*, 183(5), 1249-1263. e1223.
- Wierstra, D., & Wiering, M. (2004). *Utile distinction hidden Markov models*. Paper presented at the Proceedings of the twenty-first international conference on Machine learning.
- Wilson, R. C., Nassar, M. R., & Gold, J. I. (2010). Bayesian online learning of the hazard rate in change-point problems. *Neural computation*, 22(9), 2452-2476.
- Wilson, R. C., Nassar, M. R., & Gold, J. I. (2013). A Mixture of Delta-Rules Approximation to Bayesian Inference in Change-Point Problems. *PLoS computational biology*, 9(7), e1003150. doi:10.1371/journal.pcbi.1003150
- Wilson, R. C., Takahashi, Y. K., Schoenbaum, G., & Niv, Y. (2014). Orbitofrontal cortex as a cognitive map of task space. *Neuron*, 81(2), 267-279.
- Wood, J. N., Tierney, M., Bidwell, L. A., & Grafman, J. (2005). Neural correlates of script event knowledge: A neuropsychological study following prefrontal injury. *Cortex: A Journal Devoted to the Study of the Nervous System and Behavior*, 41(6), 796-804. doi:10.1016/S0010-9452(08)70298-3

- Yamamoto, J., & Tonegawa, S. (2017). Direct medial entorhinal cortex input to hippocampal CA1 is crucial for extended quiet awake replay. *Neuron*, 96(1), 217-227. e214.
- Yoon, K., Buice, M. A., Barry, C., Hayman, R., Burgess, N., & Fiete, I. R. (2013). Specific evidence of low-dimensional continuous attractor dynamics in grid cells. *Nature Neuroscience*, 16(8), 1077-1084.
- Yu, A. J., & Dayan, P. (2005). Uncertainty, Neuromodulation, and Attention. *Neuron*, 46(4), 681-692. doi:<https://doi.org/10.1016/j.neuron.2005.04.026>
- Zacks, J. M., Speer, N. K., Swallow, K. M., Braver, T. S., & Reynolds, J. R. (2007). Event perception: a mind-brain perspective. *Psychological Bulletin*, 133(2), 273.
- Zeithamova, D., Dominick, A. L., & Preston, A. R. (2012). Hippocampal and ventral medial prefrontal activation during retrieval-mediated learning supports novel inference. *Neuron*, 75(1), 168-179.
- Zeithamova, D., & Preston, A. R. (2010). Flexible memories: differential roles for medial temporal lobe and prefrontal cortex in cross-episode binding. *Journal of Neuroscience*, 30(44), 14676-14684.

Box inserts

Box 1. Definitions

Adaptive learning: The process of recalibrating the impact of new experience for adjustments in ongoing behavior.

Structure learning: The process of inferring the generative model that gives rise to data.

States/Latent states: the relevant variables for behavior, causal to the observations, that must be inferred (and which changes the meaning of these observations). Also called “contexts”.

Mental model: the hypothetical possible latent state(s) that people consider at a given time to make inferences about the environment

Learning rate: the degree to which immediate behavior is adjusted in accordance with a new observation (e.g., a learning rate of 1 means one would adjust their behavior entirely according to the most recent observation).

Box 2. What Happened?

Our goal is to make the best predictions possible about the next datapoint, X_{t+1} , after observing a sequence of previous data $X_{1:t}$. In principle, one might get traction on this problem by considering all possible underlying causes for the data that we have already seen. This can be done in the state framework by considering all possible sequences of state transitions underlying the observed data (denoted $S_{1:t}$). For example, after observing two data points I might consider that both came from the same state, or I might consider that these data points came from different states. Some state sequence histories may be highly consistent with previous data, whereas others may be less so. An ideal observer might attain the full predictive

distribution over future outcomes by summing across predictions made by all such state sequence histories, weighted by their likelihood of giving rise to previous data:

$$p(X_{t+1}|X_{1:t}) = \sum_{S_{1:t}} p(X_{t+1}|X_{1:t}, S_{1:t}) p(S_{1:t}|X_{1:t})$$

However, keeping track of all possible combinations of state sequence histories is intractable for most problems. The number of state sequence histories that need to be considered depends on both states and time – with n possible states the number of such histories is n^t for the most general case. Thus, practical application of this equation requires reducing the number of sequences – for example by dropping ones that fall below a certain threshold for probability (Wilson et al., 2010), storing only the K highest probability state histories (Franklin et al., 2020), or sampling from the probability distribution over state histories (Lloyd & Leslie, 2013).

Box 3. Learning over, and about states

Each state has its own set of parameters that define how observations are generated (mean, variance, etc.). We are learning at two levels when we assign data to a new state. We are both evaluating how likely the observed data fits into the current state, and we are also evaluating how likely the parameters of the state are given the data we have assigned to it. Learning about a particular state is also captured by $P(X_t|S_{1:t}, X_{1:t-1})$, if we assume that data are only used to update parameters (θ) associated with the active state. Expressing assignments of all previous data ($X_{1:t-1}$) that were assigned state S as X^S then:

$$P(X_t|S_{1:t}, X_{1:t-1}) = \int p(X_t|\theta^{S_t}, S_t) p(\theta^{S_t}|S_t, X^{S_t}) d\theta$$

Assuming that data are partitioned across states appropriately, this equation allows us to estimate probability distributions over the parameters controlling observations ($p(\theta^{S_t}|S_t, X^{S_t})$) and in turn use them to generate a probability distribution over the next observation: ($p(X_t|\theta^{S_t}, S_t)$).

Box 4. Change of states

At each point in time, we are considering a mixture of probabilities over models (e.g., economic recovery $\approx 40\%$ probable + continued recession $\approx 60\%$ probable + meteor striking earth and wiping out all humans $\approx 0.00001\%$ probable¹). The equation in Box 1 captures the idea that probabilities over many possible models are summed together to express the probability of the next data point. Clearly, this becomes very complicated and sometimes intractable, if a large number of models are being considered.

¹ Nelson, S. A. (April 27, 2018). Meteorites, Impacts, and Mass Extinction. Retrieved from https://www.tulane.edu/~sanelson/Natural_Disasters/impacts.htm

One way to simplify this task is to reduce the number of models being considered to one and to focus on the probability that the latest data point belongs to the currently active state or not. That is, is S_{t+1} the same state as S_t or not? This changepoint probability can be expressed as

$$P(S_{t+1} \neq S_t) = \frac{P(X_t|model, S_t \neq S_{t-1})H}{P(X_t|model, S_t \neq S_{t-1})H + P(X_t|model, S_t = S_{t-1})(1-H)}$$

Where $P(X_t|model, S_t \neq S_{t-1})$ is the distribution of observations predicted after a changepoint, and H is the hazard rate, which defines the transition probability associated with going from the current state to a new one. Note that the likelihood of the new observation coming from the current state is in the denominator of the equation – such that the more improbable an observation under our current assumptions about state – the more likely it is to reflect a transition to an unexpected state, in this case due to a changepoint. The same basic equation applies to the oddball and reversal tasks, with the change that $P(X_t|model, S_t \neq S_{t-1})$ reflects the predictive distribution associated with an oddball event or a reversal, and that H reflects the state transition probability associated with an oddball or reversal.