

Foundations of a temporal RL

Marc W. Howard and Zahra G. Esfahani
Boston University

Bao Le and Per B. Sederberg
University of Virginia

Recent advances in neuroscience and psychology show that the brain has access to timelines of both the past and the future. Spiking across populations of neurons in many regions of the mammalian brain maintains a robust temporal memory, a neural timeline of the recent past. Behavioral results demonstrate that people can estimate an extended temporal model of the future, suggesting that the neural timeline of the past could extend through the present into the future. This paper presents a mathematical framework for learning and expressing relationships between events in continuous time. We assume that the brain has access to a temporal memory in the form of the real Laplace transform of the recent past. Hebbian associations with a diversity of synaptic time scales are formed between the past and the present that record the temporal relationships between events. Knowing the temporal relationships between the past and the present allows one to predict relationships between the present and the future, thus constructing an extended temporal prediction for the future. Both memory for the past and the predicted future are represented as the real Laplace transform, expressed as the firing rate over populations of neurons indexed by different rate constants s . The diversity of synaptic timescales allows for a temporal record over the much larger time scale of trial history. In this framework, temporal credit assignment can be assessed *via* a Laplace temporal difference. The Laplace temporal difference compares the future that actually follows a stimulus to the future predicted just before the stimulus was observed. This computational framework makes a number of specific neurophysiological predictions and, taken together, could provide the basis for a future iteration of RL that incorporates temporal memory as a fundamental building block.

Consider the experience of listening to a familiar melody. As the song unfolds, notes feel as if they recede away from the present, an almost spatial experience. According to Husserl (1966) “points of temporal duration recede, as points of a stationary object in space recede when I ‘go away from the object.’” For a familiar melody, Husserl (1966) argues that events predicted in the future also have an analogous spatial extent, a phenomenon he referred to as *protention*. This experience is consistent with the hypothesis that the brain maintains an inner timeline extending from the distant past towards the present and from the present forwards into the future. In addition to introspection and phenomenological analysis, one can reach similar conclusions from examination of data in carefully controlled cognitive psychology experiments (Tiganj, Singh, Esfahani, & Howard, 2022).

The evolutionary utility of an extended timeline for future events is obvious. Knowing what will happen when in the future allows for selection of an appropriate action in the present. Indeed, much of computational neuroscience pre-

sumes that the fundamental goal of the cortex is to predict the future (Clark, 2013; Friston & Kiebel, 2009; Friston, 2010; Rao & Ballard, 1999; Palmer, Marre, Berry, & Bialek, 2015).

In AI, a great deal of research focuses on reinforcement learning (RL) algorithms that attempt to optimize future outcomes within a particular planning horizon (Ke et al., 2018; Dabney et al., 2020) without a temporal memory. From the perspective of psychology, RL is a natural extension of the Rescorla-Wagner model (Rescorla & Wagner, 1972) an associative model for classical conditioning (Sutton & Barto, 1981; Schultz, Dayan, & Montague, 1997; Waelti, Dickinson, & Schultz, 2001). Associative models describe connections between a pair of stimuli (or stimulus and an outcome etc) as a simple scalar value. Variables that affect the strength of an association, such as the number of pairings between stimuli, or attention, etc, all combine to affect a single scalar value. Thus, although the strength of an association can fall off with the time between stimuli, the association itself does not actually convey information about time *per se* (Gallistel, 2021a).

In RL, the goal of the Bellman equation is to estimate discounted future reward:

$$V(t) \approx \sum_{\tau} \gamma^{\tau} r(t + \tau) \quad (1)$$

but without ever explicitly estimating the future $r(t + \tau)$. Temporal difference (TD) learning requires only measurement of the reward in the present and the value of states local in time:

$$\delta(t) = r(t) + \gamma \hat{V}(t + 1) - \hat{V}(t) \quad (2)$$

Address correspondence to Marc Howard, marc777@bu.edu. The authors gratefully acknowledge support from a Google-AI Faculty Research Award and NIMH R01MH132171 (to MWH), discussions with Vijay Nambodiri, Joe McGuire, Karthik Shankar, Nao Uchida and the Uchida lab as well as careful review of previous versions of the manuscript by Nicole Howard, Hallie Botnovcan, and Aakash Sarkar.

In recent years, several authors have pursued temporal alternatives to TD learning (Ludvig, Sutton, & Kehoe, 2008; Kurth-Nelson & Redish, 2009; Momennejad & Howard, 2018; Tiganj, Gershman, Sederberg, & Howard, 2019; Tano, Dayan, & Pouget, 2020). Those models have attempted to incorporate temporal information into states (Ludvig et al., 2008) or choose a spectrum of discount rates to give appropriate behavior at a range of scales (Kurth-Nelson & Redish, 2009; Momennejad & Howard, 2018; Tano et al., 2020). To the extent that the brain can directly estimate the future, the problem solved by the Bellman equation—compute expected future reward without explicitly computing the future—is not a problem that is faced by the brain.

Within psychology, many theorists argue that classical conditioning reflects explicit storage and retrieval of temporal contingencies between stimuli (Cohen & Eichenbaum, 1993; Arcediano, Escobar, & Miller, 2005; Balsam & Gallistel, 2009; Gallistel, Craig, & Shahan, 2019). Recent neurophysiological work (Jeong et al., 2022) has shown that temporal contingency provides a better account of the firing of dopamine neurons than TDRL, highlighting the need for a neural theory for learning temporal contingency. Such a theory requires a temporal memory.

Temporal memory in the brain

There is overwhelming evidence that temporal memory is widespread throughout the mammalian brain. From examining ongoing neural activity, it is possible to decode what happened when in the past from many brain regions (King & Dehaene, 2014; Murray et al., 2017; Terada, Sakurai, Nakahara, & Fujisawa, 2017; Rossi-Pool et al., 2019; Enel, Wallis, & Rich, 2020; Cueva et al., 2020). There are at least two forms of coding for time that support this ability. So-called time cells (Pastalkova, Itskov, Amarasingham, & Buzsaki, 2008; MacDonald, Lepage, Eden, & Eichenbaum, 2011, for reviews see Eichenbaum, 2014, 2017; Tsao, Yousefzadeh, Meck, Moser, & Moser, 2022) fire in sequence following a salient stimulus. Different stimuli trigger different sequences of time cells, so that from observing which time cells are firing at any moment it is possible to decode what happened when in the past. In addition, so-called “temporal context cells” are triggered shortly after presentation of an event and then relax exponentially back to their baseline firing rate. Critically, temporal context cells have a heterogeneity of time constants (Tsao et al., 2018; Bright et al., 2020). Because different stimuli trigger different temporal context cells and because temporal context cells have a wide range of time constants, one can decode information about what happened when in the past over a wide range of time scales from a population of temporal context cells.

The properties of temporal context cells—exponential receptive fields with a wide range of time constants—are as one would expect if firing rate across the population of temporal context cells records the real Laplace transform of the past leading up to the present (Shankar & Howard, 2010; Howard et al., 2014). Let’s assume that $\mathbf{f}(t)$ is a vector describing which of several discrete events happen at time t and that $\mathbf{f}(t)$

is zero for most values of t . We can specify the past leading up to each moment t as $\mathbf{f}_t(\tau)$ (see Figure 1) with τ ranging from zero to $-\infty$ describing how far in the past an event was experienced. The goal of the temporal memory is to estimate the past $\mathbf{f}_t(\tau)$ at each time t ,

$$\mathbf{F}_t(s) = \int_0^\infty e^{-s\tau} \mathbf{f}_t(-\tau) d\tau = \mathcal{L}\{\mathbf{f}_t(-\tau)\}(s) \quad (3)$$

where we understand s to be restricted to the positive real line. The observation that $\mathbf{F}_t(s)$ is the Laplace transform of $\mathbf{f}_t(-\tau)$ establishes that it serves as a temporal memory. Time cells, with circumscribed receptive fields resemble a “direct” estimate of the past:

$$\tilde{f}_t(\tilde{\tau} < 0) = -\frac{1}{\tilde{\tau}} \int_0^\infty \Phi\left(\frac{-\tau}{\tilde{\tau}}\right) f_t(-\tau) d\tau \quad (4)$$

where $\Phi(x)$ is a unimodal function with its maximum at 1 and $\tilde{\tau}$ is defined to be negative. Time cells have properties that we would expect if the firing rate over a population of neurons records the approximate inverse Laplace transform. As $\Phi()$ becomes more and more sharp, approaching a delta function, we see that $\tilde{f}_t(\tilde{\tau})$ goes to $f_t(t + \tilde{\tau})$. The properties of time cells in the hippocampus confirm several predictions that follow from this hypothesis (Kraus, Robinson, White, Eichenbaum, & Hasselmo, 2013; Taxis et al., 2020; Cao, Bladon, Charczynski, Hasselmo, & Howard, 2022).

Eligibility traces and memory in RL

Although RL models have historically assumed Markov statistics and used the theory of Markov decision processes, the idea of a temporal memory is not at all foreign to RL. The eligibility trace (Sutton, 1988) at time t , \mathbf{e}_t , updates from time step to time step as

$$\mathbf{e}_t = \lambda \mathbf{e}_{t-1} + \mathbf{f}_t$$

where \mathbf{f}_t is the state observed at time t . It is clear that this expression results in exponential forgetting of inputs. Taking the continuum limit we find

$$\mathbf{e}_t = \sum_0^\infty \lambda^\tau \mathbf{f}_{t-\tau} \simeq \int_0^\infty e^{-s\tau} \mathbf{f}_t(\tau) d\tau \quad (5)$$

where in the last expression s is chosen as $s = -\log \lambda$. Thus the eligibility trace is an exponentially-weighted sum over past events. Comparing this last expression to Eq. 3, we find that the salient difference between the eligibility trace and Laplace transform of the past is that the eligibility trace is usually understood to have one forgetting rate λ whereas the Laplace transform requires a continuum of rate constants s .

By choosing a continuum of forgetting rates, one obtains a temporal memory extending roughly from the fastest time constant to the slowest time constant. The resolution of this temporal memory is controlled by the spacing between adjacent time constants and the degree to which the firing rates

of spiking neurons can faithfully obey Eq. 3. Critically, the resolution of this Laplace temporal memory does not depend on the properties of $\Phi()$ or whatever mechanism is used to extract information. This paper assumes the existence of a population of neurons whose firing rate $\mathbf{F}_t^-(s)$ encodes the real Laplace transform of the past leading up to time t . One of the primary goals of the paper is to develop a hypothesis that allows construction of a population $\mathbf{F}_t^+(s)$ that provides an estimate of the real Laplace transform of the future that is expected to follow time t .

Constructing neural timelines of the past and future

The goal of this paper is to write out a formalism we can take seriously to describe laboratory behavioral tasks used in psychology and neuroscience. We assume that the input to the model is a finite set of discrete symbols, x , y , etc., that are occasionally presented for an instant in continuous time. We refer to the symbol available at any particular moment as a stimulus. When a symbol is presented at time t , the stimulus is the basis vector for that stimulus for a delta function centered at t . At most times, the stimulus is zero. We assume that there are temporal relationships between some of the symbols. For convenience we assume that the time between repetitions of any symbol is much longer than the temporal relationships that are to be discovered and much longer than the longest time constant $1/s_{\min}$. This assumption allows us to imagine that experience is segmented into a series of discrete trials; each symbol can be presented at most once per trial. This assumption is not fundamental to the model but allows easy interpretation of quantities that we will derive.

The present

Let us take as input to the model a stream of inputs, $\mathbf{f}(t)$. The notation \mathbf{v} refers to a vector with each element a real number, \mathbf{v}' is a transposed vector, so that $\mathbf{u}'\mathbf{v}$ is the inner product, a scalar, and $\mathbf{u}\mathbf{v}'$ is the outer product, a matrix. We write \mathbf{f}_t for the stimulus available at time t . At instants t when no stimulus is presented, $\mathbf{f}_t = \mathbf{0}$, the vector with all entries zero. We will occasionally refer to the moment on a particular trial when x is presented, $\mathbf{f}_t = \mathbf{x}$ as t_x . We ignore similarity between symbols, so that $\mathbf{y}'\mathbf{x} = \delta_{y,x}$. If symbol x was predicted to occur at time t , but was not observed, we will occasionally write $\mathbf{f}_t = \tilde{\mathbf{x}}$ to describe the observation of a failed prediction for symbol x . One may imagine the basis vectors \mathbf{x} , \mathbf{y} , $\tilde{\mathbf{x}}$ etc as one-hot vectors without changing any of the results in this paper.

We write $\mathbf{f}_t(\tau)$ to describe the true past that led up to time t , where τ runs from 0^- , corresponding to the moment of the past closest to the present backwards to $-\infty$, corresponding to the distant past. Whereas \mathbf{f}_t is the stimulus available in the present at time t , $\mathbf{f}_t(\tau < 0)$ is the timeline that led up to time t . Under the assumption that every symbol is presented at most once per trial, each component of $\mathbf{f}_t(\tau < 0)$ is either a delta function at some particular τ or zero everywhere.

Neural manifolds for the past and the future

We estimate both the past and the future as functions over neural manifolds. Each manifold is a population of processing elements—neurons—each of which is indexed by a position in a coordinate space. The coordinates describing the neurons are continuous and locally Euclidean. At each moment, each neuron is mapped onto scalar value correspond to its firing rate over a macroscopic period of time on the order of at least tens of ms. We propose that the past and the future are represented by separate manifolds that interact with one another.

The representations for both the past and the future each utilize two connected manifolds. We refer to one kind of manifold, indexed by an effectively continuous variable s , as a Laplace space. The other kind of manifold, indexed by an effectively continuous variable τ^* , is referred to as an inverse space. The representations of the past follow previous work in theoretical neuroscience (Shankar & Howard, 2013; Howard et al., 2014), psychology (Howard, Shankar, Aue, & Criss, 2015), and neuroscience (Bright et al., 2020; Cao et al., 2022).

Laplace spaces for remembered past and predicted future. The Laplace space corresponding to the past, which we write as $\mathbf{F}_t^-(s)$ encodes the Laplace transform of $\mathbf{f}_t(\tau)$, the past leading up to time t :

$$\mathbf{F}_t^-(s) = \mathcal{L}\{\mathbf{f}_t(\tau < 0)\}(s) \quad (6)$$

We restrict s to real values on the positive line (but see Aghajani, Kreiman, & Fried, 2022). The Laplace space corresponding to the future, which we write as $\mathbf{F}_t^+(s)$ is an attempt to estimate the Laplace transform of the future, $\mathcal{L}\{\mathbf{f}_t(\tau > 0)\}(s)$. Many neurons tile the s axis continuously for each symbol. One may imagine that each symbol, rather than being represented by a single neuron as in a one-hot vector, is represented by a line of neurons representing the history, or future, of that symbol. The index s assigned to a neuron corresponds to the inverse of its functional time constant. Thus, there is a natural mapping between $1/s$ and τ within both the past and the future. By convention, s is positive for both the past and the future so that $\mathbf{F}_t^-(s)$ is the Laplace transform of $\mathbf{f}_t(-\tau)$ for $\tau < 0$ whereas $\mathbf{F}_t^+(s)$ is the Laplace transform of $\mathbf{f}_t(\tau)$ for $\tau > 0$.

Although s is effectively continuous, this does not require that neurons sample s evenly. Following previous work in psychology (e.g., Chater & Brown, 2008; Piantadosi, 2016; Howard & Shankar, 2018), neuroscience (Guo, Huson, Mascosko, & Regehr, 2021; Cao et al., 2022), and theoretical neuroscience (Lindeberg & Fagerström, 1996; Shankar & Howard, 2013), we assume that s is sampled on a logarithmic scale. Let n be the neuron number, starting from the largest value of s_{\max} nearest $\tau = 0$ and extending out from the present. We obtain a logarithmic scale by choosing $ds/dn = -s$.

Updating Laplace spaces in real time. Suppose that we have arranged for one particular component of $\mathbf{F}_t^-(s)$ or

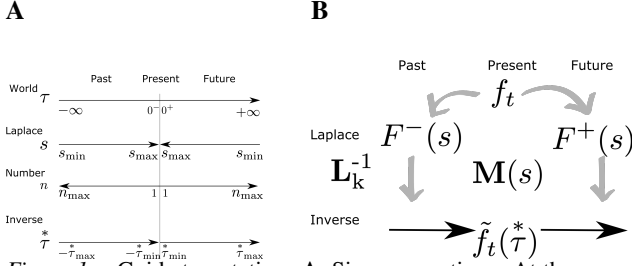


Figure 1. Guide to notation. **A.** Sign conventions. At the present moment t , objective time τ runs from $-\infty$ to ∞ . $\tau = 0$ corresponds to time t . The real Laplace domain variable s runs from 0^+ to $+\infty$ for both past and future, approximated as s_{\min} and s_{\max} . The units of s are t^{-1} ; the values corresponding to different points of the timeline are shown in the same vertical alignment. Cell number for Laplace and inverse spaces n are aligned with one another. The variable *taustar* describes position along the inverse spaces. It is in register with τ and derived from τ . **B.** The stimulus available in the present, \mathbf{f}_t provides input to two sets of neural manifolds. One set of neural manifolds represents the past; the other estimates the future. $\mathbf{M}(s)$ stores temporal relationships between events.

$F_t^+(s)$ to hold the Laplace transform of one particular symbol, which we write as $f_t(\tau)$. Suppose further that $f_t(\tau)$ is zero in the neighborhood of $\tau = 0$. Consider how this component, which we write as $F^-(s)$ or $F^+(s)$, should update as time passes. Let us pick some minimal increment of time δt on the order of, say, 100 ms. At time $t + \delta t$, information in $f_t(\tau < 0)$ recedes further away from the present, so that $F_{t+\delta t}^-(s) = \mathcal{L}\{f_t(\tau + \delta t)\}$. In contrast, at time $t + \delta t$, information in $f_t(\tau > 0)$ comes closer to the present, so that $F_{t+\delta t}^+(s) = \mathcal{L}\{f_t(\tau - \delta t)\}$. More generally, suppose that $F_t(s)$ is the Laplace transform of a function over some variable x , $F_t(s) = \mathcal{L}\{f_t(x)\}(s)$. Defining $\alpha \equiv \delta x / \delta t$, we can update $F_t(s)$ as

$$F_{t+\delta t}(s) = \mathcal{L}\{\mathcal{T}_{\alpha(\delta t)}f_t(x)\}(s) = e^{-s\alpha(\delta t)}F_t(s) \quad (7)$$

where \mathcal{T} is the translation operator, $\mathcal{T}_a f(x) = f(x+a)$ and we have used the expression for the Laplace transform of translated functions. Equation 7 describes a recipe for updating both $F_t^\pm(s)$ with α_\pm in the absence of new input. Using the sign convention developed here, we fix $\alpha_- = 1$ for $F^-(s)$ and fix $\alpha_+ = -1$ for $F^+(s)$. It is possible to incorporate changes into the rate of flow of subjective time by letting α_\pm change in register, such that $\alpha_+(t) = -\alpha_-(t)$ for all t . The expression in Eq. 7 holds more generally and can be used to update Laplace transforms over many continuous variables of interest for cognitive neuroscience (Howard et al., 2014; Howard, Luzzardo, & Tiganj, 2018; Howard & Hasselmo, 2020).

We are in a position to explain how $\mathbf{F}_t^-(s)$ comes to represent the Laplace transform of $\mathbf{f}_t(\tau < 0)$; a discussion of how $\mathbf{F}_t^+(s)$ comes to estimate the future requires more development and will be postponed. When a symbol is presented at time t , it enters timeline of the past at $\tau = 0^-$. So, incorporating the input at time t into the past at time $t + \delta t$ we

have

$$\begin{aligned} \mathbf{F}_{t+\delta t}^-(s) &= e^{-s(\delta t)} [\mathbf{F}_t^-(s) + \mathcal{L}\{\delta(0^-)\mathbf{f}_t\}] \\ &= e^{-s(\delta t)}\mathbf{F}_t^-(s) + e^{-s(\delta t)}\mathbf{f}_t. \end{aligned} \quad (8)$$

At time $t + \delta t$, the input from time t is encoded as the Laplace transform of that symbol a time δt in the past. At each subsequent time step, an additional factor of $e^{-s\delta t}$ accumulates. As time passes, the input from time t is always stored as Laplace transform of a delta function at the appropriate place on the timeline. Because this is true for all stimuli that enter $\mathbf{F}^-(s)$, we conclude that $\mathbf{F}_t^-(s)$ encodes the Laplace transform of the past $\mathbf{f}_t(\tau < 0)$.

The middle panel of Figure 2 illustrates the profile of activity over \mathbf{F}_t^- and \mathbf{F}_t^+ , shown as a function of cell number n , resulting from the Laplace transform of a delta function at various moments in time. In the middle panel, the axis for the past is reversed to allow appreciation of the relationship between past time $\tau < 0$ and F^- . Note that the Laplace transform of a delta function has a characteristic shape as a function of cell number that merely translates as time passes. Note that the magnitude of the translation of $F^\pm[n]$ depends on the value of τ_o . It can be shown that for a delta function $F_{t+\delta t}^\pm[n] = F_t^\pm[n + \delta n]$ with $\delta n = \alpha_\pm \frac{\delta t}{\tau_o}$. This can be appreciated by noting that the distances between successive lines in the middle panel of Fig. 2 are not constant despite the fact that they correspond to the same time displacement. Whereas δn goes down as time passes for $F^-[n]$, δn increases with the passage of time for $F^+[n]$.

There are implementational challenges to building a neural circuit that obeys Eq. 7; these challenges are especially serious when $\alpha < 0$, which requires activation to grow exponentially. These challenges would be mitigated by a neural circuit that obeys an equivalent PDE obtained by differentiating Eq. 7 with respect to cell number n . The use of a PDE allows error to be distributed over many neurons and would allow neurons that have zero activation to grow if their neighbors are nonzero. Moreover, if one could literally implement the PDE this would preserve linearity. A disadvantage of a PDE is that it may require careful tuning of a neural circuit. If one were willing to restrict the representation of each symbol to the Laplace transform of a delta function at a single point in time, it would be straightforward to implement a continuous attractor network (Khona & Fiete, 2021) to allow the ‘‘edge’’ in the Laplace transform as a function of n to translate appropriately.

Inverse spaces for remembered past and predicted future. The bottom panel of Figure 2 shows a graphical depiction of the inverse space for the past and the future during the interval between presentation of x and y . The inverse spaces approximate the past, $\tilde{f}_t^*(\tau < 0)$, and the future, $\tilde{f}_t^*(\tau > 0)$, on a log scale. Neurons in the inverse space have circumscribed receptive fields in time, like time cells in the hippocampus (Pastalkova et al., 2008; MacDonald et al., 2011). As the delta function corresponding to the time of x recedes into the past, the corresponding bump of activity in $\mathbf{x}'\tilde{f}_t^*(\tau < 0)$ also moves, keeping its shape but moving more and more slowly

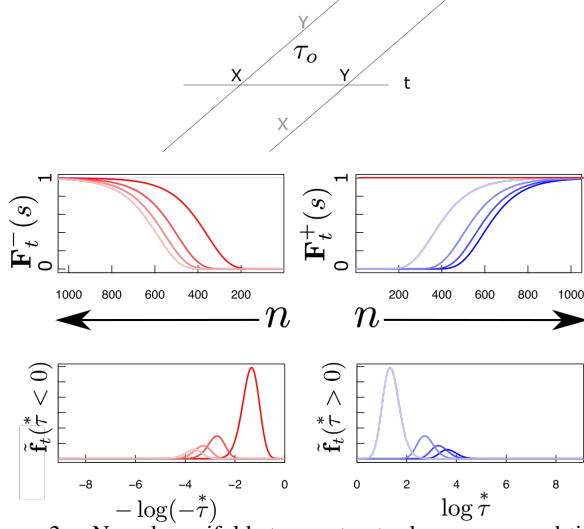


Figure 2. Neural manifolds to construct a log compressed timeline of the past and the future. Top: A temporal relationship exists between x and y such that y always follows x after a delay of τ_o seconds. Consider how the internal timeline ought to behave after x is presented at $t = 0$. At time t , the past should include x t seconds in the past and y $\tau_o - t$ seconds in the future. Samples of the timeline at evenly-spaced moments between zero and τ_o . Earlier moments closer to $t = 0$ are darker and later moments closer to $t = \tau_o$ are lighter. Red lines are neurons coding for x , blue lines are neurons coding for y . Middle: Laplace spaces for the past (left) and future (right) shown as a function of cell number n ; Bottom: inverse spaces, constructed using the Post approximation, for the past (left) and future (right) shown as a function of log time. Exactly at time $t = 0$, x is available a time 0^+ in the future (dark horizontal red line, middle right). Similarly, exactly at $t = \tau_o$, y is available a time 0^- in the past (light horizontal blue line, middle left).

as x recedes further and further into the past. In the future, the delta function corresponding to the predicted time of y should start a time τ_o in the future and come closer to the present as time passes. As the prediction for y approaches the present, the corresponding bump of activity in $\mathbf{y}^* \hat{\mathbf{f}}_t^+(\tau^* > 0)$ keeps its shape but the speed of the bump accelerates rather than slowing with the passage of time.

Previous papers have made use of the Post approximation to implement the inverse transform. This is not neurally reasonable (Gosmann, 2018); the Post approximation is difficult to implement even in artificial neural networks (e.g., Tano et al., 2020; Jacques, Tiganj, Howard, & Sederberg, 2021). A more robust approach would be a continuous attractor network (for a review see Khona & Fiete, 2021) that takes input as the derivative of F with respect to n . The width of the bump in $\hat{\mathbf{f}}$ would depend on internal connections between neurons in $\hat{\mathbf{f}}$ and global inhibition would stabilize the activity over $\hat{\mathbf{f}}$. In this case, moving the bump in different directions, corresponding to $\alpha > 0$ and $\alpha < 0$ is analogous to moving the bump of activity in, say, a ring attractor for the head direction system, in different directions.

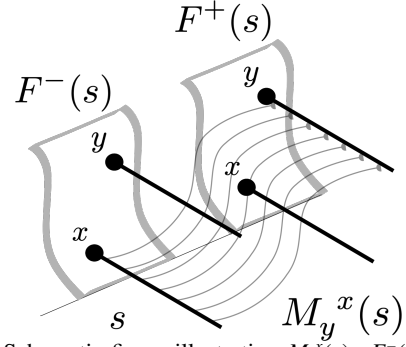


Figure 3. Schematic figure illustrating $M_y^x(s)$. $F^-(s)$ and $F^+(s)$ components for all the possible symbols, here shown schematically as sheets. Two symbols x and y are shown in both $F^-(s)$ and $F^+(s)$. Each symbol is associated with a population of neurons spanning a continuous set of s values, shown as the heavy lines in this cartoon. $\mathbf{M}(s)$ describes the connections between each symbol in $F^-(s)$ to each symbol in $F^+(s)$ for each value of s . The curved lines $M_y^x(s)$ illustrate the set of weights connecting units corresponding to x in F^- to units corresponding to y in F^+ . Connections exist only between units with the same values of s . The strength of the connections in $M_y^x(s)$ vary as a function of s in a way that reflects the pairwise history between x and y .

Predicting the future from the past

The previous subsection describes how to evolve the Laplace manifold for the past. We could use the same approach to evolve the Laplace manifold for the future during periods when no symbol is experienced if we could initialize the representation of the future appropriately. This will be accomplished *via* learned temporal relationships between the past and the future. For present we only consider simple pairwise relationships between symbols.

The moment a nonzero stimulus \mathbf{f}_t is experienced, we assume it is available to both \mathbf{F}^- and \mathbf{F}^+ , triggering a number of operations which presumably occur sequentially within a small window of time on the order of perhaps 100 ms. First, the present stimulus updates a prediction for the future *via* a set of connections \mathbf{M} organized by s . Then these connections are updated by associating the past to the present. Finally the present stimulus is added to the representation of the past. For ease of exposition we will first focus on describing the connections between the past and the future.

We write $\mathbf{M}(s)$ for a set of connections that associates the Laplace transform of the past to the Laplace transform of the future (Fig. 3). For any particular value s_o , $\mathbf{M}(s_o)$ is a matrix describing connections from each symbol in $\mathbf{F}^-(s_o)$ to each symbol in $\mathbf{F}^+(s_o)$. For each pair of symbols, say x and y , we write $M_y^x(s_o)$ for the strength of the connection *from* the cell corresponding to x with $s = s_o$ in F^- to the cell corresponding to y in F^+ with $s = s_o$. $\mathbf{M}(s)$ does not include connections between neurons with different values of s . On occasion it will be useful to think of the set of connections between a pair of symbols over all values of s , which we write as $M_y^x(s)$. Similarly, we write $\mathbf{M}^y(s)$ for the set of connections *from* y in F^- to all stimuli in F^+ over all values of

s . We write $\mathbf{M}_y(s)$ for the set of connections to y in F^+ from all symbols and all values of s .

When a particular stimulus, let's say y , is presented, the connections to and from that stimulus in $\mathbf{M}(s)$ are updated. The connections from y in the past are updated as

$$\mathbf{M}^y(s) \rightarrow \rho \mathbf{M}^y(s) \quad (9)$$

That is, the connections from y in F^- to every other stimulus for each value of s , are all scaled down by a value ρ . Later we will consider the implications of a continuous spectrum of ρ values; for now let us just treat ρ as a fixed parameter restricted to be between zero and one. When y is presented, it momentarily becomes available at the ‘‘rearward part’’ of the future. In much the same way that the present enters the past (Eq. 8) at $\tau = 0^-$, we also assume that the present is also available momentarily in the future at $\tau = 0^+$. The connections to y in the future are updated as

$$\mathbf{M}_y(s) \rightarrow \mathbf{M}_y(s) + (1 - \rho) \mathbf{F}_t^-(s) \quad (10)$$

We can understand Eq. 10 as a Hebbian association between the units in $\mathbf{F}^-(s)$, whose current activation is given by $\mathbf{F}_t^-(s)$ and the units in the future $\mathbf{F}^+(s)$ corresponding to the present stimulus y (see Fig. 4). More generally, we can understand this learning rule as strengthening connections from the past $\mathbf{F}_t^-(s)$ to the rearward part of the future, $\mathcal{L}\{\delta(0^+) \mathbf{f}_t\}(s) = e^{-s0^+} \mathbf{f}_t$. Because the second term is the product of two Laplace transforms, it can also be understood as the Laplace transform of a convolution, here, the convolution of the present with the past.¹ Convolution has long been used as an associative operation in mathematical psychology (Murdock, 1982; Jones & Mewhort, 2007), neural networks (Plate, 1995; Eliasmith, 2013; Blouw, Solodkin, Thagard, & Eliasmith, 2016), and computational neuroscience (Steinberg & Sompolinsky, 2022).

$\mathbf{M}(s)$ stores pairwise temporal relationships. To understand the properties of $\mathbf{M}(s)$, let us assume that the model as described thus far learns in a world containing two stimuli, x and y for many trials. Let us assume that x is presented on each trial. If y is presented on a given trial it appears precisely τ_o seconds after x . The probability y is presented on each trial is $P(y|x)$. In this simple situation we can restrict our attention to $M_y^x(s)$, the weights connecting the neurons coding for x in the past to the neurons coding for y in the future.

From examination of Eqs. 9 and 10, we see that after each trial $M_y^x(s)$ is multiplied by ρ when x was presented. For trials on which y was also presented, $(1 - \rho)e^{-s\tau_o}$ is added to $M_y^x(s)$. Writing $h[i]$ as an indicator variable for the history of presentations of y on the trial i steps in the past we find

$$M_y^x(s) = (1 - \rho) e^{-s\tau_o} \sum_i \rho^i h[i] \quad (11)$$

Note that if $P(y|x) = 1$, then after an infinitely long series of trials $\sum_i h[i] \rho^i = \frac{1}{1-\rho}$ and $M_y^x(s) = e^{-s\tau_o}$ for all choices of ρ . Following similar logic, if we relax the assumption that

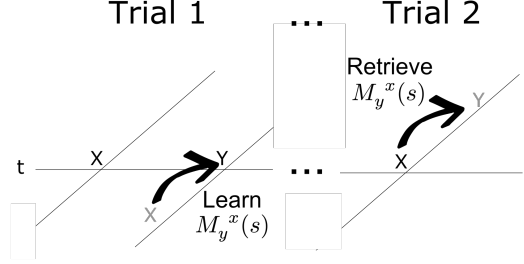


Figure 4. Learning and expressing pairwise associations with $\mathbf{M}(s)$. The horizontal line is time; the diagonal lines indicate the internal timeline at the moments they intersect. Memory for the past is below the horizontal line; prediction of the future is above. When x is presented for the first time, it predicts nothing. When y is presented, the past contains a memory for x in the past. When y is presented, $M_y^x(s)$ stores the temporal relationship between x in the past and y in the present—the rearward part of the future. In addition to storing learned relationships, connections from each item decay each time it was presented (not shown). When x is repeated much later in time, the stored connections in $M_y^x(s)$ retrieve a prediction of y in the future.

$P(y|x) = 1$ and take the limit as ρ goes to 1, we find that $M_y^x(s) = P(y|x) e^{-s\tau_o}$.

Now let us relax the assumption that the time lag between x and y always takes the same value. Let the lag be a random variable τ_{xy} subject to the constraint that τ_{xy} is always > 0 . This is not a fundamental restriction; if τ_{xy} changed sign, those observations would contribute to $M_x^y(s)$ instead of $M_y^x(s)$. Now, again taking the limit as $\rho \rightarrow 1$, we find

$$M_y^x(s) = P(y|x) E[e^{-s\tau_{xy}}] = P(y|x) \mathcal{L}\{\tau_{xy}\}(s) \quad (12)$$

where we have used the definition for the Laplace transform of a random variable, again with the understanding that we restrict s to be real and positive.

Equation 12 illustrates several important properties of $\mathbf{M}(s)$. First, we can see that $M_y^x(s)$ provides complete information about the distribution of temporal lags averaged over history. This can be further appreciated by noting that the Laplace transform of the random variable on the right hand side is the moment generating function of $-\tau_{xy} = \tau_{yx}$. Keeping the computation in the Laplace domain means that there is no blur introduced by going into the inverse space as in previous attempts to build a model for predicting the future (Momennejad & Howard, 2018; Tiganj et al., 2019; Goh, Ursekar, & Howard, 2022). Second, because $\mathcal{L}\{\tau_{xy}\}(s = 0) = 1$ as long as the expectation of τ_{xy} is finite, we can write $M_y^x(s) = M_y^x(s = 0) \hat{M}_y^x(s)$ where $M_y^x(s = 0)$ is just $P(y|x)$. This allows us to cleanly decompose information about what will happen in the future, stored in $M_y^x(s = 0)$, from information about when those events will happen, stored in $\hat{M}_y^x(s)$

¹ Because of the sign conventions adopted here, $F_t^-(s)$ is the Laplace transform of $f_t(-\tau)$ whereas $F_t^+(s)$ is the transform of $f_t(\tau)$. Viewed in this light it is more precise to think of Eq. 10 as learning the Laplace transform of the cross-correlation between the present and the past.

More generally we can express $\mathbf{M}(s)$ as $\mathbf{M}(s) = \mathbf{M}_{\text{what}} \cdot \hat{\mathbf{M}}(s)$ where \cdot indicates pointwise rather than matrix multiplication.

Continuum of ρ and memory for trial history. Before moving on we briefly note the implications of understanding ρ as a continuous variable. Treating ρ as continuous, Eq. 11, which describes the situation where τ_{xy} is equal to τ_0 on each trial, can be rewritten as

$$M_y^x(\rho, s) = (1 - \rho)e^{-s\tau_0} \mathcal{Z}\{h[i]\}(\rho^{-1})$$

where $\mathcal{Z}\{z\}$ is the Z-transform, the discrete analog of the Laplace transform (Ogata, 1970).

Although the notation is a bit more unwieldy, allowing τ_{xy} to vary across trials we see that the trial history of timing is also retained by $\mathbf{M}(\rho, s)$. Writing the delay between x and y on the trial i steps in the past as $\tau[i]$, and $H[i](s) \equiv h[i]e^{-s\tau[i]}$

$$M_y^x(\rho, s) = (1 - \rho)\mathcal{Z}\{H[i](s)\}(\rho^{-1}). \quad (13)$$

Because the Z-transform is in principle invertible, information about the entire trial history has been retained by virtue of having a continuum of forgetting rates ρ . Figure 5 illustrates the ability to extract the trial history including timing information of events that follow x from $\mathbf{M}(\rho, s)\mathbf{x}$. This illustrates a remarkable property of Laplace-based temporal memory. Although each synaptic matrix with a specific value of ρ_0 forgets exponentially with rate $-\log \rho_0$, the set of matrices with a continuum of ρ retains information about the entire trial history. Of course, in practice there must be some numerical imprecision in the biological instantiation of $\mathbf{M}(\rho, s)$. In principle however, a continuum of forgetting rates ρ means that the past is not forgotten. Rather the past, here as a function of trial history, has been written to the continuous values of ρ .

Updating the future

Let us return to the problem of generating a prediction of the immediate future. We again restrict our attention to the limit as ρ goes to 1 and assume the system has experienced a very long sequence of trials with the same underlying statistics. Moreover, we assume for the present that only pairwise relationships are important, so we can neglect the temporal credit assignment problem, and construct the Laplace transform of the future that predicted solely on the basis of the present stimulus.

There are two problems that need to be resolved to write an analog of Eq. 8 for $\mathbf{F}_{t+\delta t}^+(s)$. First, we can only use use Eq. 7 to update $\mathbf{F}_t^+(s)$ if $\mathbf{F}_t^+(s)$ is already the Laplace transform of a predicted future; we must create a circumstance that makes that true. Second, we need to address the situation where a prediction reaches the present. Because of the discontinuity at $\tau = 0$ special considerations are necessary to allow the time of a stimulus to pass from the future to the past.

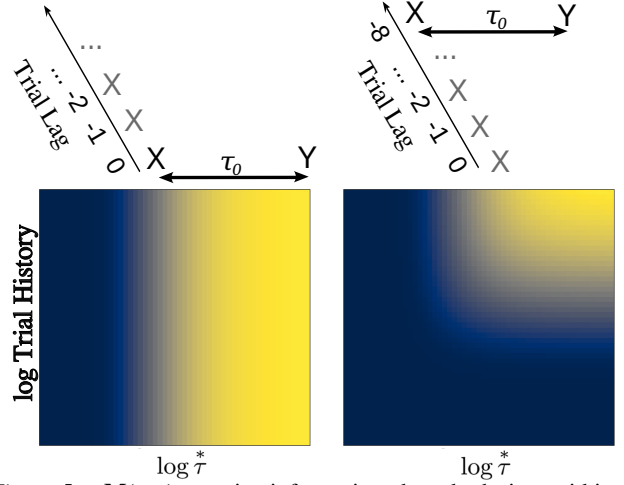


Figure 5. $\mathbf{M}(\rho, s)$ contains information about both time within a trial and trial history. Left: Consider a single pairing of x and y on the most recent trial. The heatmap shows the degree to which y is cued by x by $\frac{y^T \mathbf{M}(\rho, s) \mathbf{x}}{1 - \rho}$ projected onto log time. The profile as a function of $\log \tau$ is identical to the profile for future time in Figure 2. If the pairing between x and y had a longer delay, the edge would be further to the right. Right: The single pairing of x and y is followed by an additional series of trials on which x was presented by itself. Now there is an edge in both trial history and time within trial. Additional trials with only x would push this edge further towards the top of the graph. Additional trials with x and y paired would be added to this plot with a time delay that reflects the timing of the pairing.

Predicting the future with the present. Equation 12 indicates that the weights in $M_y^x(s)$ record the future occurrences of y given that x occurs in the present. $M_y^x(s)$ captures both the probability that y will follow x as well as the distribution of temporal delays at which y is expected to occur. This information is encoded as a probability times the Laplace transform of a random variable. If we only need to consider x in predicting the future, then $M_y^x(s)$ is precisely how we would like to initialize the future prediction for y in $\mathbf{F}_t^+(s)$ after x is presented (Fig. 4).

We probe $\mathbf{M}(s)$ with the “immediate past.” When x is presented it enters $\mathbf{F}_t^-(s)$ as $\mathcal{L}\{\delta(0^-)\mathbf{x}\}(s)$. Multiplying $\mathbf{M}(s)$ from the right with the immediate past, yields a prediction for the future.

$$\mathbf{M}(s)e^{-s0}\mathbf{x} = \mathbf{M}(s)\mathbf{x} = P(y|x)\mathcal{L}\{\tau_{xy}\}\mathbf{y} \quad (14)$$

More generally, the input to the future at time t should be given by $\mathbf{M}(s)\mathcal{L}\{\delta(0^-)\mathbf{f}_t\}$. For concision we write this as $\mathbf{M}(s)\mathbf{f}_t$. Because the past stored in $\mathbf{M}(s)$ was a probability times the Laplace transform of a probability distribution, so too will the future recovered in this way.

Continuity of the predicted future through $\tau = 0$. The neural representation described here approximates a continuous timeline by stitching together separate Laplace neural man-

ifolds for the past and the future. With the passage of time, information in the future moves ever closer to the present. As time passes and a prediction reaches the present, this discontinuity must be addressed.

We can detect predictions that have reached the present by examining $\mathbf{F}_t^+(s = \infty)$, which only rises from zero when $\tau \rightarrow 0$. In practice, we would use s_{\max} which should be on the order of $(\delta t)^{-1}$. If the future that is being represented is the Laplace transform of a delta function, then we can simply take components for which $\mathbf{F}_t^+(s_{\max}) > 0$ to zero for all s at the next time step. More generally, if the future that is represented is not simply a delta function, the linearity of the Laplace transform allows us to subtract $\mathbf{F}_t^+(s = \infty)$ from all s values without affecting the evolution at subsequent time points.

If a prediction reaches the present and is observed, then no further action is needed. If a prediction reaches the present, but is not observed, we can trigger an observation of a “not stimulus”, written e.g., \tilde{x} to describe the observation of a failed prediction for a stimulus x . Although we won’t pursue it here, one could allow “not stimuli” to be predicted by stimuli and to predict other stimuli, allowing for the model to provide appropriate predictions for a relatively complex set of contingencies.

Evolving $\mathbf{F}_{t+\delta t}^+(s)$. Integrating these two additional factors allows us to write a general expression for evolving $\mathbf{F}_t^+(s)$ to $\mathbf{F}_{t+\delta t}^+(s)$.

$$\mathbf{F}_{t+\delta t}^+(s) = e^{-s\delta t}\mathbf{F}_t^+(s) - \mathbf{F}_t^+(s = \infty) + \mathbf{M}(s)\mathbf{f}_t. \quad (15)$$

In simple situations where it is sufficient to know pairwise associations between stimuli separated in time, this provides a complete model for constructing a timeline of the future.

Credit assignment and Laplace temporal difference learning

Consider an experiment. In the control condition of this hypothetical experiment, the participant is presented with y followed by z at a delay of 5 s for N trials. In the experimental condition, during an initial phase of training, x is followed by z at a delay of 10 s for some number of trials. After this initial training, the participant is presented with x followed by y at a delay of 5 s and then z 5 s after y for N trials. The number of pairings between y and z , and the delays between them, are identical in the two conditions so that $M_z^y(s)$ would be the same in the two conditions. However, whereas y is “solely responsible” for z in the control condition, x is also capable of predicting z in the experimental condition. We might expect the “credit” that y receives for predicting z to be less in the experimental condition and indeed, analogous effects are observed in temporal blocking experiments (e.g., Amundson & Miller, 2008).

Traditional RL makes sense of this phenomenon, and blocking more generally, by hypothesizing that plasticity at the time z is presented depends on how well it was predicted. However, in the present framework, in the experimental condition, z would be predicted an appropriate distance in the

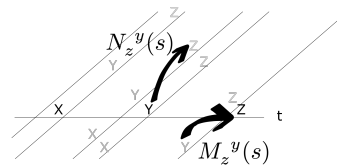


Figure 6. Whereas $\mathbf{M}(s)$ records the Laplace transform of the future that follows each item, $\mathbf{N}(s)$ records the expectation of the Laplace transform of the future predicted prior to that item. Consider a scenario with three consecutive items x , y , and z so that z is predicted prior to the presentation of y . $M_z^y(s)$ learns the connection between y in the past and z in the present. $N_z^y(s)$ learns the connection between y in the present and the prediction for z available just before y was presented. Comparing $M_z^y(s)$ to $N_z^y(s)$ allows one to estimate how much y is responsible for presentation of z .

future at the moment x is presented. The two conditions thus differ in the degree to which z is predicted at the moment that y is presented. The basic strategy pursued here is to use conditions at the time of y to assess how much credit it should get for the occurrence of z , including the time of its presentation (Figure 6).

We assume that there is a prediction $\mathbf{F}_t^+(s)$ available at all times, although this prediction can be zero. The “Laplace temporal difference” measure compares the true future that follows each stimulus, stored in $\mathbf{M}(s)$, to the future predicted just before each stimulus is presented. A second set of connections, $\mathbf{N}(s)$, records the average future available just before each item was observed. That is, for each stimulus y , $\mathbf{N}^y(s)$ averages $\mathbf{F}_t^+(s)$ observed at moments t such that $\mathbf{f}_t = y$. The same strategy was used in Goh et al. (2022). We will see that this Laplace temporal difference is sensitive to the amount of information about both the identity and timing of future events. As such it provides a neurally-reasonable mechanism to estimate temporal contingency (Balsam & Gallistel, 2009; Gallistel, 2021b; Jeong et al., 2022).

Estimating prediction independent of the present via $\mathbf{N}(s)$

If y is presented at time t , then $\mathbf{N}(s)$ is updated as

$$\mathbf{N}^y(\rho, s) \rightarrow \rho\mathbf{N}^y(\rho, s) + (1 - \rho)\mathbf{F}_t^+(s) \quad (16)$$

Following Eq. 13 we see that if ρ is a continuous variable, then $\mathbf{N}^y(\rho, s)$ is the Z-transform of the trial history of the predicted future available just prior to the moment y was presented. Note that, like $\mathbf{M}(\rho = 1, s)$, $\mathbf{N}(\rho = 1, s)$ can be decomposed into components corresponding to information about what symbols are predicted to occur and when. Defining $\mathbf{N}_{\text{what}} = \mathbf{N}(s = 0)$, we can write $\mathbf{N}(s) = \mathbf{N}_{\text{what}} \cdot \hat{\mathbf{N}}(s)$ much like we did for $\mathbf{M}(s)$.

To illustrate the properties of $\mathbf{M}(\rho, s)$ and $\mathbf{N}(\rho, s)$, let us restrict our attention to cases where at most three symbols x , y and z are presented in order on each trial. Let us refer to the time lags between symbols as random variables τ_{xy} , τ_{yz} ; on trials where all three symbols are observed $\tau_{xz} = \tau_{xy} + \tau_{yz}$ by assumption. We assume that the distributions are chosen

such that the relative times of presentation do not overlap. We denote the probabilities of each symbol occurring on a trial such that $P(z|y)$ gives the conditional probability that z is observed on a trial given that y is also observed on that trial.

Estimating three point correlation functions from $\mathbf{M}(\rho)$ and $\mathbf{N}(\rho)$

A great deal of information can be extracted from the trial history encoded in $\mathbf{M}(\rho, s)$ and $\mathbf{N}(\rho, s)$. On a particular trial after y has been presented, we would like to compute the probability that z will be presented and at what time. \mathbf{M} contains the two-point probability distribution of y and z . It would be preferable to predict the occurrence of z using the three-point probability distribution, taking into account the occurrence and timing of both x and y .

Because $\mathbf{M}(\rho, s)$ and $\mathbf{N}(\rho, s)$ contain information about the paired trial history, in principle we can extract information about the three-point correlation function. For instance, if z only occurs on trials on which both x and y are presented, then we should observe a positive correlation between the trial history encoded in $M_z^y(\rho, s = 0)$ and the trial history of predictions encoded in $N_z^y(\rho, s = 0)$. After all, if x is not presented on a particular trial, then it cannot predict z and the prediction written into $N_z^y(\rho, s = 0)$ on that trial will be zero. Similarly, one can imagine that the joint timing of the presentations of x and y predicts the timing of z . For instance, a positive correlation between τ_{xy} and τ_{yz} would be observed if x causes the future occurrence of both y and z *via* some process whose rate varies across trials. If τ_{xy} is longer than average on a particular trial, then z will be predicted earlier than average on that trial. In this way, a positive correlation between τ_{xy} and τ_{yz} would manifest as a negative correlation across trials between the time at which z is observed and the time at which z is predicted. In principle this information can be extracted from $\hat{\mathbf{M}}_z^y(\rho, s)$ and $\hat{\mathbf{N}}_z^y(\rho, s)$.

Associative contingency via Laplace temporal difference

A general model for predicting the future taking into account the past and the present, as well as the trial history encoded in $\mathbf{M}(\rho, s)$ and $\mathbf{N}(\rho, s)$ is beyond the scope of this paper. We will content ourselves, taking the limit as $\rho \rightarrow 1$, with a simple measure to estimate contingency between pairs of symbols. If y has no causal effect on the future that follows it, then $\mathbf{M}^y(s) = \mathbf{N}^y(s)$. This holds for what and when information taken separately. Let us define a simple measure of contingency between symbols, which we refer to as the Laplace temporal difference, as

$$\text{LTD}(s) \equiv -\log \frac{\mathbf{M}(s)}{\mathbf{N}(s)} = -\log \frac{\mathbf{M}_{\text{what}}}{\mathbf{N}_{\text{what}}} - \log \frac{\hat{\mathbf{M}}(s)}{\hat{\mathbf{N}}(s)} \quad (17)$$

where we understand the fractions of matrices on the right-hand-side to imply pointwise division. The ratio $\frac{\mathbf{M}(s)}{\mathbf{N}(s)}$ can be understood as an attempt to deconvolve the predictions that

precede each symbol from the future that follows it. When a symbol does not predict the future occurrence of symbols that follow it, the deconvolution yields the Laplace transform of delta function at zero. The Laplace transform of $\delta(0) = 1$ for all s , so that $\text{LTD}(s)$ is zero for all s under these circumstances.

We will work out the properties of $\text{LTD}(s)$ in our simple case with three symbols to illustrate that it functions as a reasonable measure of contingency. We are interested in the contingency between y and z given the context provided by x .

From Eq. 12, we know that $M_z^y(s) = P(z|y)\mathcal{L}\{\tau_{yz}\}$. Because $\tau_{yz} = \tau_{xz} - \tau_{xy}$ by assumption, it is also true that $M_z^y(s) = P(z|y)\mathcal{L}\{\tau_{xz} - \tau_{xy}\}$. To compute $N_z^y(s)$ we first need to determine the prediction from x to z on each trial y is presented. Then we need to average that prediction over trials. Because $N_z^y(s)$ is only updated on trials where y is presented, x can only contribute to $N_z^y(s)$ when x is also presented, which occurs with probability $P(x|y)$. On each trial, right after x is presented the prediction for z is $M_z^x(s) = P(z|x)\mathcal{L}\{\tau_{xz}\}$. On each trial, at the time y is observed this contribution from x has evolved to $e^{s(t_y - t_x)}P(z|x)\mathcal{L}\{\tau_{xz}\}$, where $t_y - t_x$ is the value of the random variable τ_{xy} observed on that particular trial. Averaging over trials on which x is presented gives the Laplace transform of the convolution of τ_{xz} and $-\tau_{xy}$. Multiplying by the probability that x is presented at all we find

$$N_z^y(s) = P(z|x)P(x|y)\mathcal{L}\{\tau_{xz} * (-\tau_{xy})\}(s). \quad (18)$$

Now,

$$-\log \frac{\mathbf{z}'\mathbf{M}_{\text{what}}\mathbf{y}}{\mathbf{z}'\mathbf{N}_{\text{what}}\mathbf{y}} = -\log \frac{P(z|y)}{P(z|x)P(x|y)} \quad (19)$$

Equation 19 shows that the Laplace temporal difference identifies differences between future outcomes that are expected to follow the present stimulus—the numerator—and future outcomes that would be predicted by previous stimuli—the denominator. Although Eq. 19 is written for a single preceding symbol, this result obviously generalizes over mutually exclusive symbols that could precede y , $\sum_x P(z|x)P(x|y)$.

Considering temporal contingency, we find

$$-\log \frac{\hat{M}_z^y(s)}{\hat{N}_z^y(s)} = -\log \frac{\mathcal{L}\{\tau_{xz} - \tau_{xy}\}(s)}{\mathcal{L}\{\tau_{xz} * (-\tau_{xy})\}(s)} \quad (20)$$

Recall that the expectation of the sum of two independent random variables is equal to their convolution. The ratio in Equation 20 compares the Laplace transform of the sum of two random variables to the Laplace transform of their convolution. If τ_{xz} and $-\tau_{xy}$ are independent, then the ratio is 1 and the right hand side is zero for all s . This demonstrates that $\text{LTD}(s)$ is sensitive to the temporal contingency between y and z .

Neural predictions

Regions as widely separated as the cerebellum (Wagner & Luo, 2020; De Zeeuw, Lisberger, & Raymond, 2021),

striatum (e.g., van der Meer & Redish, 2011), PFC (e.g., Rainer, Rao, & Miller, 1999; Ning, Bladon, & Hasselmo, 2022), OFC (e.g., Namboodiri et al., 2019; Schoenbaum, Chiba, & Gallagher, 1998; Young & Shapiro, 2011), hippocampus (Ferbinteanu & Shapiro, 2003; Duvelle, Grieves, & van der Meer, 2022) and thalamus (Komura et al., 2001) contain active representations that code for the future. One can find evidence of predictive signals extending over long periods of time that modulate firing in primary visual cortex (Gavornik & Bear, 2014; Kim, Homann, Tank, & Berry, 2019; Homann, Koay, Chen, Tank, & Berry, 2022; Yu et al., 2022). Prediction apparently involves a substantial proportion of the brain. Coordinating activity and plasticity over such a wide region would require careful synchronization (Hasselmo, Bodelón, & Wyble, 2002; Hamid, Frank, & Moore, 2021). The timescale of this synchronization, presumably on the order of 100 ms, fixes δt , places a bound on the fastest timescales $1/s$ that can be observed, and operationalizes the duration of the “present.”

Given the widespread nature of predictive signal, we will not attempt to map specific equations onto specific brain circuits. Rather we will illustrate the observable properties implied by these equations with an eye towards facilitating future empirical work. The predictions fall into two categories. One set of predictions describes properties of ongoing firing of neurons participating in Laplace and inverse spaces. Another set of predictions are a direct consequence of the properties of learned weights. We also briefly discuss the model in this paper in the context of recent empirical work on the computational basis of the dopamine signal (Jeong et al., 2022).

Active firing neurons

This paper proposes the existence of neural manifolds to code for the identity and time of *future* events. The prediction is that there should be two related manifolds, one implementing the Laplace space and one implementing the inverse space. Previous neuroscientific work has shown evidence for Laplace and inverse spaces for a timeline for the past. The properties of the proposed neural manifolds for future time can be understood by analogy to the neural manifolds for the past.

Single-cell properties of neurons coding for the past. So-called temporal context cells observed in the entorhinal cortex (Tsao et al., 2018; Bright et al., 2020) are triggered by a particular event and then relax exponentially back to baseline firing with a variety of time constants. The firing of temporal context cells is as one would expect for a population coding $\mathbf{F}^-(s)$. So-called time cells observed in the hippocampus (Pastalkova et al., 2008; MacDonald et al., 2011; Taxis et al., 2020; Shahbaba et al., 2022; Shikano, Ikegaya, & Sasaki, 2021; Schonhaut, Aghajian, Kahana, & Fried, 2022) and many other brain regions (e.g., Tiganj, Cromer, Roy, Miller, & Howard, 2018; Tiganj, Kim, Jung, & Howard, 2017; Mello, Soares, & Paton, 2015; Bakhurin et al., 2017; Akhlaghpour et al., 2016; Jin, Fujii, & Graybiel, 2009) fire sequentially as events recede into the past, as one would

expect from neurons participating in $\tilde{f}^*(\tau < 0)$. Time cells are consistent with qualitative and quantitative predictions of $\tilde{f}^*(\tau < 0)$, including the conjecture that time constants are distributed along a logarithmic scale (Cao et al., 2022).

Single-cell and population-level properties of neurons coding for the past and the future. In situations where the future can be predicted, $\mathbf{F}^+(s)$ and $\tilde{f}^*(\tau > 0)$ should behave as mirror images of the corresponding representations of the past. Figure 7A illustrates the firing of cells coding for a stimulus remembered in the past (left) and predicted in the future (right). Neurons participating in the Laplace space, sorted on their values of s , are shown in the top; neurons participating in the inverse space, sorted on their values of τ are shown on the bottom.

The firing of neurons constituting the Laplace space shows a characteristic shape when plotted as a function of time in this simple experiment. Neurons coding for the past are triggered shortly after presentation of the stimulus and then relax exponentially with a variety of rates. Neurons coding for the future ramp up, peaking as the predicted time of occurrence grows closer. The ramps have different characteristic time constants. Different populations are triggered by the presentation of different symbols (not shown) so that the identity of the remembered and predicted symbols as well as their timing can be decoded from populations coding $\mathbf{F}^-(s)$ and $\mathbf{F}^+(s)$. The largest value of $1/s$ in the figure is chosen to be a bit longer than the delay in the experiment, resulting in a subset of neurons that appear to fire more or less constantly throughout the delay (Enel et al., 2020).

The firing of neurons constituting the inverse space also shows a characteristic shape when plotted as a function of time in this simple experiment. Neurons tile the delay, with more cells firing early in the interval with more narrow receptive fields. The logarithmic compression of n results in a characteristic “backwards J” shape for the past and a mirror image “J” shape for the future. Again, different populations would code for different stimuli in the past and in the future (not shown) so that the identity of the remembered and predicted stimuli and their time from the present could be decoded from a population coding $\tilde{f}^*(\tau)$. Figure 7B shows firing that would be expected for a population that includes cells coding for the same stimulus, say γ , both in the past and the future around the time of a predicted occurrence of that symbol.

Plausible anatomical locations for an internal future timeline. This computational hypothesis should be evaluated with carefully planned analyses. However, the published literature shows evidence that is at least roughly in line with the hypothesis of neural manifolds for future time. Firing that ramps systematically upward in anticipation of important outcomes including planned movements has been observed in (at least) mPFC (Henke et al., 2021), ALM (Inagaki, Inagaki, Romani, & Svoboda, 2018), and thalamus (Komura et al., 2001). Komura et al. (2001) showed evidence for ramping firing in the thalamus that codes for outcomes in

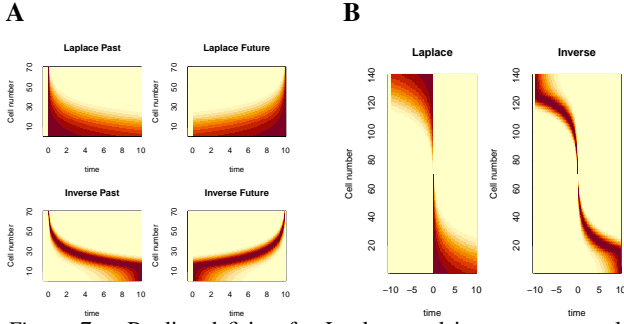


Figure 7. Predicted firing for Laplace and inverse spaces plotted as heatmaps. **A.** Consider an experiment in which x precedes y separated by 10 s. The top row shows firing as a function of time for cells in the Laplace space for the past (left) and the future (right). Note that the cells in $F_t^-(s)$ peak at time zero and then decay exponentially. In contrast cells in $F_t^+(s)$ peak at 10 s and ramp up exponentially. The bottom row shows firing as a function of time for cells in the Inverse space. **B.** Consider an experiment in which y is predicted to occur at time zero and then recedes into the past. Cells coding for both past and future are recorded together and sorted on the average time at which they fire. Left: For Laplace spaces, neurons in $F_t^+(s)$ are sorted to the top of the figure and neurons $F_t^-(s)$ are sorted to the bottom of the figure. Right: Inverse spaces show similar properties but give rise to a characteristic “pinwheel” shape.

a Pavlovian conditioning experiment. Preliminary evidence from secondary analyses suggest that there is a heterogeneity of time constants in ALM (Inagaki et al., 2018) and mPFC (Henke et al., 2021).

There is also circumstantial neurophysiological evidence for sequential firing leading to predicted future events as predicted by $\tilde{f}(\tau > 0)$. Granule cells in cerebellum appear to fire in sequence in the time leading up to an anticipated reward (Wagner, Kim, Savall, Schnitzer, & Luo, 2017; Wagner & Luo, 2020). OFC may be another good candidate region to look for “future time cells.” OFC has long been argued by many authors to code for the identity of predicted outcomes (Hikosaka & Watanabe, 2000; Schoenbaum & Roesch, 2005; Mainen & Kepecs, 2009). More recently Enel et al. (2020) showed sequential activation in OFC during a task in which it was possible to predict the value of a reward that was delayed for several seconds. Finally, it should be noted that the properties of $\tilde{f}(\tau > 0)$ are a temporal analog of spatial “distance-to-goal” cells observed in spatial navigation studies (Sarel, Finkelstein, Las, & Ulanovsky, 2017; Gauthier & Tank, 2018).

Predictions from weight matrix $\mathbf{M}(\rho, s)$

Properties of weights due to s . Consider an experiment in which different symbols, denoted cs1, cs2, etc, precede an outcome r by a delay τ_o . The value of τ_o changes across the different symbols (Figure 8A). Ignoring ρ for the moment, the strength of the connections from each cs to r depend on the value of τ_o for that stimulus and the value of s for each synapse: $e^{-s\tau_o}$. When a particular cs is presented at time t ,

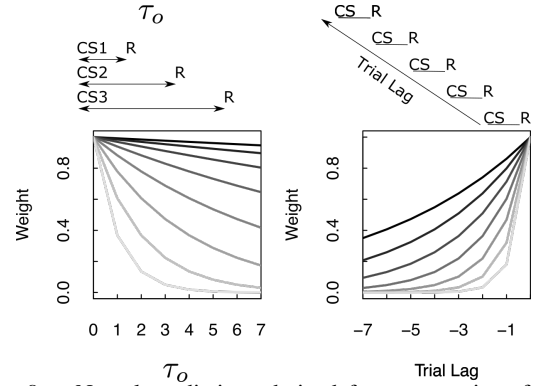


Figure 8. Neural predictions derived from properties of $\mathbf{M}(s)$. **Left.** Plot of the magnitude of the entry in $\mathbf{M}_r(\rho = 1, s)$ connecting each of the conditioned stimuli cs to the outcome r as a function of the τ_o corresponding to that cs. Different lines correspond to entries with different values of s . Weights corresponding to different values of s show exponential discounting as τ_o is manipulated, with a variety of discount rates. **Right.** Plot of the magnitude of $\mathbf{M}(\rho, s = 0)$ associated with a single pairing of cs and r a certain number of trials in the past. Different lines show the results for different values of ρ . For clarity, these curves have been normalized such that they have the same value at trial lag zero.

the amount of information that flows along each synapse is $e^{-s\tau_o}$ and the pulse of input to $\mathbf{F}_{t+\delta t}^+(s) - \mathbf{F}_t^+(s)$ corresponding to the outcome is $e^{-s\tau_o}$.

Thus, considering each connection as a function of τ_o , firing should go down exponentially as a function of τ_o with a rate constant that depends on the value of s . This pattern of results aligns well with results presented at SfN in 2022 (Masset, Malik, Kim, Bech Vilaseca, & Uchida, 2022). The experiment was constructed much as described above. Masset et al. (2022) measured the firing of dopamine neurons to different stimuli predicted reward delivery at different delays. It has long been known that firing of dopamine neurons, averaged over neurons, around the time of the conditioned stimulus goes down with delay (Fiorillo, Newsome, & Schultz, 2008). This study showed that there was a heterogeneity of exponential decay rates in the firing of dopamine neurons in this paradigm, much as illustrated in Fig. 8A. In the context of TDRL, this finding is also consistent with a continuous spectrum of exponential discount rates (Momennejad & Howard, 2018; Tano et al., 2020).

Properties of weights due to ρ . A continuum of forgetting rates ρ predicts a range of trial history effects. Figure 8B shows the weights in $\mathbf{M}(\rho)$ over past trials that result from different values of ρ . This is simply ρ^i where i is the trial recency with values normalized such that the weight at the most recent trial is 1. The weights $\mathbf{M}(\rho)$ record the trial history of reinforcement. The weights $\mathbf{N}(\rho)$ record the trial history of predicted outcomes. The difference between $\mathbf{M}(\rho)$ and $\mathbf{N}(\rho)$ as a function records the history of prediction violations. Many papers show dependence on previous trial outcomes in response to a cue stimulus in learning and decision-making experiments (Bernacchia, Seo, Lee, & Wang, 2011;

Morcos & Harvey, 2016; Scott et al., 2017; Akrami, Kopec, Diamond, & Brody, 2018; Hattori, Danskin, Babic, Mlynaryk, & Komiyama, 2019; Hattori & Komiyama, 2022). These studies show history-dependent effects in a wide range of brain regions and often show a continuous spectrum of decay rates within a brain region (see especially Bernacchia et al., 2011). Notably, distributions of time constants for trial history effects cooccur with distributions of ongoing activity in multiple brain regions (Spitman, Seo, Lee, & Soltani, 2020).

Dopamine and learning

The connection between TDRL and neuroscience related to dopamine has been one of the great triumphs of computational neuroscience (Schultz et al., 1997). The standard account is that the firing of dopamine neurons signals reward prediction error (RPE) which drives plasticity. Despite its remarkable success at predicting the findings of many behavioral and neurophysiological experiments, the RPE account has been under increasing strain over recent years. The standard account did not predict the existence of a number of striking effects, including increasing dopamine firing during delay under uncertainty (Fiorillo, Tobler, & Schultz, 2003), dopamine ramps in spatial experiments (Howe, Tierney, Sandberg, Phillips, & Graybiel, 2013), dopamine waves (Hamid et al., 2021), and heterogeneity of dopamine responses across neurons and brain regions (Dabney et al., 2020; Masset et al., 2022; W. Wei, Mohebi, & Berke, 2021). Recently Jeong et al. (2022) reported the results of several experiments that flatly contradict the standard model. These experiments were proposed to evaluate an alternative hypothesis for dopamine firing in the brain.

Jeong et al. (2022) propose that dopamine signals whether the current stimulus is a cause of reward. The model developed there, referred to as ANCCR, assesses the contingency between a stimulus and outcomes. $LTD(s)$ measures the contingency—temporal and otherwise—between a symbol and possible outcomes. Both ANCCR and the framework developed in this paper are inspired by a similar critique of Rescorla-Wagner theory and TDRL (Gallistel, 2021b). In order to make a complete account of the experiments in the Jeong et al. (2022) paper, the current framework would have to be elaborated in several ways. In order to keep the calculations simple here we have assumed that each symbol can only occur once per trial. Consider the $\mathbf{M}(s)$ and $\mathbf{N}(s)$ that would result if this assumption were relaxed. If we define $p_y^x(\tau)$ as the probability that we will observe y in a small interval around $t + \tau$ given that we observed x at time t , we get

$$M_z^y(s) = \int e^{-s\tau} p_z^y(\tau) d\tau = \mathcal{L}\{p_z^y\}(s).$$

Assuming we can pretend that it is acceptable to ignore the overlap between the distributions we would find

$$\begin{aligned} N_z^y(s) &= \sum_x \int e^{-s\tau} (p_y^x \# p_z^x)(\tau) d\tau \\ &= \mathcal{L}\{p_y^x \# p_z^x\}(s). \end{aligned}$$

Where $\#$ signifies cross-correlation. This is closely analogous to Eqs. 19 and 20.

The expression for $\mathbf{M}(s)$ ends up giving expected number of observations of z that would follow y out to a timescale on the order of $1/s$. $N_z^y(s)$ gives the number of observations of z that were expected prior to observation of y . It is thus possible to construct a measure of prospective contingency like that used in the Jeong et al. (2022) model. The current framework does not require one to specify an intrinsic timescale of association *a priori*.

Discussion

This paper takes a phenomenological approach to computational neuroscience. The strategy is to write down equations that, if the brain could somehow obey them, would be consistent with a wide range of observed cognitive and neural phenomena. The phenomenological equations make concrete predictions that can be evaluated with cognitive and neurophysiological experiments. To the extent the predictions hold, the question of how the brain manages to obey these phenomenological equations could then become a separate subject of inquiry. The phenomenological equations require a number of capabilities of neural circuits, both at the level of synapses and in terms of ongoing neural activity. We make those explicit here.

Circuit assumptions for synaptic weights

The connections $\mathbf{M}(\rho, s)$ and $\mathbf{N}(\rho, s)$ require that the brain uses continuous variables, s and ρ , to organize connections between many neurons, most likely spanning multiple brain regions. For the phenomenological equations to be viable, these continuous variables should be deeply embedded in the functional architecture of the brain. For instance, in order to invert the integral transforms, it is necessary to compute a derivative over these continuous variables. This suggests a gradient in these continuous variables should be anatomically identifiable. Conceivably anatomical gradients in gene expression and functional architecture (e.g., Phillips et al., 2019; Guo et al., 2021; Roy, Zhang, Halassa, & Feng, 2022) could generate anatomical gradients in s and/or ρ . Perhaps part of the function of traveling waves of activity such as theta oscillations (Lubenov & Siapas, 2009; Patel, Fujisawa, Berényi, Royer, & Buzsáki, 2012; Zhang & Jacobs, 2015) or dopamine waves (Hamid, Frank, & Moore, 2019) is to make anatomical gradients salient.

$LTD(s)$ is not intended as a literal description of a neural computation the brain could use to assess contingency between stimuli. However, the idea of assessing contingency by comparing the future that immediately precedes a stimulus to the future that follows it *should* be taken seriously. If brain oscillations are important in synchronizing the flow of information proposed in this framework, then perhaps information from $\mathbf{M}(s)$ and $\mathbf{N}(s)$ are available at different phases of an oscillation. Perhaps this consistent relationship temporal combined with spike timing dependent plasticity (Bi & Poo, 1998; Dan & Poo, 2004) somehow facilitates computation of the contingency between the present and the future.

Circuit assumptions for ongoing activity

At the neural level, this framework assumes the existence of circuits that can maintain activity of a Laplace neural manifold over time. There is evidence that the brain has found some solution to this problem, at least for past time and at least in the entorhinal cortex (Tsao et al., 2018; Bright et al., 2020). Exponential growth of firing, as proposed by Eq. 15 seems on its face to be a computationally risky proposition. However, this proposal does create testable predictions. Moreover, firing rates that increase monotonically in time are widely observed. For instance border cells as an animal approaches a barrier (Solstad, Boccara, Kropff, Moser, & Moser, 2008) and evidence accumulation cells (Roitman & Shadlen, 2002) both increase monotonically. If this monotonic increase in firing reflects movement of an edge along a Laplace neural manifold, the characteristic time scale of the increase should be heterogeneous across neurons. If the brain has access to a circuit with paired α 's, it could reuse this circuit to construct cognitive models for spatial navigation (Howard et al., 2014), evidence accumulation (Howard et al., 2018), and perhaps cognitive computation more broadly (Howard & Hasselmo, 2020). Consistent with this hypothesis, monotonic cells in spatial navigation and evidence accumulation—border cells and evidence accumulation cells—have sequential analogues (Wilson & McNaughton, 1993; Morcos & Harvey, 2016; Koay, Charles, Thiberge, Brody, & Tank, 2022) as one would expect if they reflect a Laplace space that is coupled with an inverse space.

Perhaps part of the solution to implementing these equations in the brain is to restrict the kinds of functions that can be represented over the Laplace manifold. Perhaps a continuous attractor network that can maintain and evolve the Laplace transform of a single delta function per basis vector would be straightforward to construct. In this case, each component of $\mathbf{F}_t^-(s)$ and $\mathbf{F}_t^+(s)$ would be at any moment the Laplace transform of a delta function; $\mathbf{M}(s)$ and $\mathbf{N}(s)$ would still be able to store distributions over multiple presentations. In this case perhaps $\mathbf{F}_{t+\delta t}^+(s)$ could update by sampling from the distribution expressed by $\mathbf{M}(s)\mathbf{f}_t$. Perhaps predictions are updated in the more general case by sampling from a superposition of the previous prediction $\mathbf{F}_t^+(s)$ and the future that follows the current item $\mathbf{M}s\mathbf{f}_t$ masked by $\mathbf{N}(s)\mathbf{f}_t$.

What about control?

RL algorithms have been successful in AI application because of their ability to learn policies to control actions in the absence of explicit supervision (Watkins & Dayan, 1992; Sutton & Barto, 2018). The current framework does not include a deep connection to control theory. It is conceivable that the current framework could be integrated into existing deep network approaches to RL. Perhaps it is possible to learn an embedding that maps continuous features onto symbols then control actions using existing methods from deep RL. Another possibility is to develop a control theory in the Laplace domain. Indeed, this is how control theory problems are typically solved analytically (Ogata, 1970). Indeed, there

is some evidence that control systems in neurobiology, for instance gaze stabilization in the oculomotor system, make use of multiple time constants over several orders of magnitude (Miri, Bhasin, Aksay, Tank, & Goldman, 2022). Perhaps a continuous set of time constants, as required for the Laplace neural manifolds used here, may enable brains to make use of diversity enabled sweet spots (Nakahira, Liu, Sejnowski, & Doyle, 2021).

It should be possible to extend the current framework to multiple dimensions beyond time, including real space and abstract spaces (Howard et al., 2014, 2018). Properties of the Laplace domain enable data-independent operators that enable efficient computation (Howard & Hasselmo, 2020). For instance, given that a state of a Laplace neural manifold is Laplace transform of a function, we can construct the Laplace transform of the translated function (Eq. 7, see also (Shankar, Singh, & Howard, 2016)). Critically, the translation operator is independent of the function to be translated. Restricting our attention to Laplace transforms of delta functions, we can construct the sum or difference using convolution and cross correlation respectively (Howard et al., 2015; Howard & Hasselmo, 2020). The binary operators for addition and subtraction also do not need to be learned. Perhaps the control theory that governs behavior is analogous to generic spatial navigation in a continuous space.

Scale-covariance as a design goal

Because the s values are sampled along a logarithmic scale, all of the quantities in this paper are scale-covariant. Rescaling time, taking $\tau_{xz} \rightarrow a\tau_{xz}$, $\tau_{xy} \rightarrow a\tau_{xy}$, etc, simply takes $s \rightarrow s/a$. Because the s values are chosen in a geometric series, rescaling time simply translates along the n axis. All the components of the model, \mathbf{F}^- , \mathbf{F}^+ , \mathbf{M} , \mathbf{N} , and $\text{LTD}(s)$, all use the same kind of logarithmic scale for time. This means that rescaling time simply translates all the components of the entire model, up to a scaling factor. All of the components of the model are thus time scale-covariant, responding to rescaling time with a translation over cell number. Thus any measure that integrates over n (and is not subject to edge effects) is scale-invariant.

Empirically, there is not a characteristic time scale to associative learning (Balsam & Gallistel, 2009; Gershman, 2022); any model that requires choice of a time scale for learning to proceed is thus incorrect. Logarithmic time scales are observed neurally (Cao et al., 2022; Guo et al., 2021). Logarithmic time scales can be understood as a commitment to a world with power law statistics (X.-X. Wei & Stocker, 2012; Piantadosi, 2016) or as an attempt to function in many different environments without a strong prior on the time scales it will encounter (Howard & Shankar, 2018).

Recent work has shown that the use of logarithmic time scales enables scale-invariant CNNs for vision (Jansson & Lindeberg, 2021) and audition (Jacques, Tiganj, Sarkar, Howard, & Sederberg, 2022). For instance, Jacques et al. (2022) trained deep CNNs to categorize spoken digits. When tested on digits presented at very different speeds than the training examples (imagine someone saying the word

“seven” stretched over four seconds), the deep CNN with a logarithmic time axis generalized perfectly. Rescaling time translates the neural representation at each layer; convolution is translation equivariant; including a maxpool operation over the convolutional layer renders the entire CNN translation-invariant. Time is not only important in speech perception (e.g., Lerner, Honey, Katkov, & Hasson, 2014) but vision as well (Russ, Koyano, Day-Cooney, Perwez, & Leopold, 2022) suggesting that these ideas can be incorporated into a wide range of sensory systems.

References

- Aghajan, Z. M., Kreiman, G., & Fried, I. (2022). Periodic time cells in human entorhinal cortex. *bioRxiv*.
- Akhlaghpour, H., Wiskerke, J., Choi, J. Y., Taliaferro, J. P., Au, J., & Witten, I. (2016). Dissociated sequential activity and stimulus encoding in the dorsomedial striatum during spatial working memory. *eLife*, 5, e19507.
- Akrami, A., Kopec, C. D., Diamond, M. E., & Brody, C. D. (2018). Posterior parietal cortex represents sensory history and mediates its effects on behaviour. *Nature*, 554(7692), 368–372.
- Amundson, J. C., & Miller, R. R. (2008). Cs-us temporal relations in blocking. *Learning & Behavior*, 36(2), 92–103.
- Arceiano, F., Escobar, M., & Miller, R. R. (2005). Bidirectional associations in humans and rats. *Journal of Experimental Psychology: Animal Behavior Processes*, 31(3), 301–18.
- Bakhurin, K. I., Goudar, V., Shobe, J. L., Claar, L. D., Buonomano, D. V., & Masmanidis, S. C. (2017). Differential encoding of time by prefrontal and striatal network dynamics. *Journal of Neuroscience*, 37(4), 854–870.
- Balsam, P. D., & Gallistel, C. R. (2009). Temporal maps and informativeness in associative learning. *Trends in Neuroscience*, 32(2), 73–78.
- Bernacchia, A., Seo, H., Lee, D., & Wang, X. J. (2011). A reservoir of time constants for memory traces in cortical neurons. *Nature Neuroscience*, 14(3), 366–72.
- Bi, G.-q., & Poo, M.-m. (1998). Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *Journal of Neuroscience*, 18(24), 10464–10472.
- Blouw, P., Solodkin, E., Thagard, P., & Eliasmith, C. (2016). Concepts as semantic pointers: A framework and computational model. *Cognitive science*, 40(5), 1128–1162.
- Bright, I. M., Meister, M. L. R., Cruzado, N. A., Tiganj, Z., Bufalo, E. A., & Howard, M. W. (2020). A temporal record of the past with a spectrum of time constants in the monkey entorhinal cortex. *Proceedings of the National Academy of Sciences*, 117, 20274–20283.
- Cao, R., Bladon, J. H., Charczynski, S. J., Hasselmo, M., & Howard, M. (2022). Internally generated time in the rodent hippocampus is logarithmically compressed. *eLife*, <https://doi.org/10.7554/eLife.75353>.
- Chater, N., & Brown, G. D. A. (2008). From universal laws of cognition to specific cognitive models. *Cognitive Science*, 32(1), 36–67. doi: 10.1080/03640210701801941
- Clark, A. (2013). Whatever next? predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(03), 181–204.
- Cohen, N. J., & Eichenbaum, H. (1993). *Memory, amnesia, and the hippocampal system*. Cambridge, MA: The MIT Press.
- Cueva, C. J., Saez, A., Marcos, E., Genovesio, A., Jazayeri, M., Romo, R., ... Fusi, S. (2020). Low-dimensional dynamics for working memory and time encoding. *Proceedings of the National Academy of Sciences*, 117(37), 23021–23032.
- Dabney, W., Kurth-Nelson, Z., Uchida, N., Starkweather, C. K., Hassabis, D., Munos, R., & Botvinick, M. (2020). A distributional code for value in dopamine-based reinforcement learning. *Nature*, 1–5.
- Dan, Y., & Poo, M. M. (2004). Spike timing-dependent plasticity of neural circuits. *Neuron*, 44(1), 23–30.
- De Zeeuw, C. I., Lisberger, S. G., & Raymond, J. L. (2021). Diversity and dynamism in the cerebellum. *Nature neuroscience*, 24(2), 160–167.
- Duvelle, É., Grieves, R. M., & van der Meer, M. (2022). Temporal context and latent state inference in the hippocampal “splitter” signal. *PsyArXiv*, <https://doi.org/10.31234/osf.io/9z4wr>.
- Eichenbaum, H. (2014). Time cells in the hippocampus: a new dimension for mapping memories. *Nature Reviews Neuroscience*, 15(11), 732–44. doi: 10.1038/nrn3827
- Eichenbaum, H. (2017). Memory: organization and control. *Annual review of psychology*, 68, 19–45.
- Eliasmith, C. (2013). *How to build a brain: A neural architecture for biological cognition*. Oxford University Press.
- Enel, P., Wallis, J. D., & Rich, E. L. (2020). Stable and dynamic representations of value in the prefrontal cortex. *Elife*, 9, e54313.
- Ferbinteanu, J., & Shapiro, M. L. (2003). Prospective and retrospective memory coding in the hippocampus. *Neuron*, 40(6), 1227–39.
- Fiorillo, C. D., Newsome, W. T., & Schultz, W. (2008). The temporal precision of reward prediction in dopamine neurons. *Nature Neuroscience*. doi: 10.1038/nn.2159
- Fiorillo, C. D., Tobler, P. N., & Schultz, W. (2003). Discrete coding of reward probability and uncertainty by dopamine neurons. *Science*, 299(5614), 1898–902. doi: 10.1126/science.1077349
- Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11, 127–138.
- Friston, K., & Kiebel, S. (2009). Predictive coding under the free-energy principle. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1521), 1211–21. doi: 10.1098/rstb.2008.0300
- Gallistel, C. (2021a). The physical basis of memory. *Cognition*, 213, 104533.
- Gallistel, C. (2021b). Robert Rescorla: Time, information and contingency. *Revista de Historia de la Psicología*, 42(1), 7–21.
- Gallistel, C., Craig, A. R., & Shahan, T. A. (2019). Contingency, contiguity, and causality in conditioning: Applying information theory and weber’s law to the assignment of credit problem. *Psychological review*, 126(5), 761.
- Gauthier, J. L., & Tank, D. W. (2018). A dedicated population for reward coding in the hippocampus. *Neuron*, 99, 179–193.e7.
- Gavornik, J. P., & Bear, M. F. (2014). Learned spatiotemporal sequence recognition and prediction in primary visual cortex. *Nature neuroscience*, 17(5), 732–737.
- Gershman, S. J. (2022). The molecular memory code and synaptic plasticity: a synthesis. *arXiv preprint arXiv:2209.04923*.
- Goh, W. Z., Ursekar, V., & Howard, M. W. (2022). Predicting the future with a scale-invariant temporal memory for the past. *Neural Computation*, 34(642–685).
- Gosmann, J. (2018). *An integrated model of context, short-term, and long-term memory* (Unpublished doctoral dissertation). University of Waterloo.
- Guo, C., Huson, V., Macosko, E. Z., & Regehr, W. G. (2021). Graded heterogeneity of metabotropic signaling underlies a con-

- tinuum of cell-intrinsic temporal responses in unipolar brush cells. *Nature Communications*, 12(1), 1–12.
- Hamid, A. A., Frank, M. J., & Moore, C. I. (2019). Dopamine waves as a mechanism for spatiotemporal credit assignment. *bioRxiv*, 729640.
- Hamid, A. A., Frank, M. J., & Moore, C. I. (2021). Wave-like dopamine dynamics as a mechanism for spatiotemporal credit assignment. *Cell*, 184(10), 2733–2749.
- Hasselmo, M. E., Bodelón, C., & Wyble, B. P. (2002). A proposed function for hippocampal theta rhythm: Separate phases of encoding and retrieval enhance reversal of prior learning. *Neural Computation*, 14, 793–817.
- Hattori, R., Danskin, B., Babic, Z., Mlynaryk, N., & Komiyama, T. (2019). Area-specificity and plasticity of history-dependent value coding during learning. *Cell*, 177(7), 1858–1872.
- Hattori, R., & Komiyama, T. (2022). Context-dependent persistency as a coding mechanism for robust and widely distributed value coding. *Neuron*, 110(3), 502–515.
- Henke, J., Bunk, D., von Werder, D., Häusler, S., Flanagan, V. L., & Thurley, K. (2021). Distributed coding of duration in rodent prefrontal cortex during time reproduction. *Elife*, 10, e71612.
- Hikosaka, K., & Watanabe, M. (2000). Delay activity of orbital and lateral prefrontal neurons of the monkey varying with different rewards. *Cerebral cortex*, 10(3), 263–271.
- Homann, J., Koay, S. A., Chen, K. S., Tank, D. W., & Berry, M. J. (2022). Novel stimuli evoke excess activity in the mouse primary visual cortex. *Proceedings of the National Academy of Sciences*, 119(5), e2108882119.
- Howard, M. W., & Hasselmo, M. E. (2020). Cognitive computation using neural representations of time and space in the Laplace domain. *arXiv preprint arXiv:2003.11668*.
- Howard, M. W., Luzardo, A., & Tiganj, Z. (2018). Evidence accumulation in a Laplace decision space. *Computational Brain and Behavior*, 1, 237–251.
- Howard, M. W., MacDonald, C. J., Tiganj, Z., Shankar, K. H., Du, Q., Hasselmo, M. E., & Eichenbaum, H. (2014). A unified mathematical framework for coding time, space, and sequences in the hippocampal region. *Journal of Neuroscience*, 34(13), 4692–707. doi: 10.1523/JNEUROSCI.5808-12.2014
- Howard, M. W., & Shankar, K. H. (2018). Neural scaling laws for an uncertain world. *Psychological Review*, 125, 47–58. doi: 10.1037/rev0000081
- Howard, M. W., Shankar, K. H., Aue, W., & Criss, A. H. (2015). A distributed representation of internal time. *Psychological Review*, 122(1), 24–53.
- Howe, M. W., Tierney, P. L., Sandberg, S. G., Phillips, P. E., & Graybiel, A. M. (2013). Prolonged dopamine signalling in striatum signals proximity and value of distant rewards. *Nature*, 500(7464), 575–579.
- Husserl, E. (1966). *The phenomenology of internal time-consciousness*. Bloomington, IN: Indiana University Press.
- Inagaki, H. K., Inagaki, M., Romani, S., & Svoboda, K. (2018). Low-dimensional and monotonic preparatory activity in mouse anterior lateral motor cortex. *Journal of Neuroscience*, 38(17), 4163–4185.
- Jacques, B. G., Tiganj, Z., Howard, M. W., & Sederberg, P. B. (2021). DeepSITH: Efficient learning via decomposition of what and when across time scales. In M. Ranzato, A. Beygelzimer, P. Liang, J. Vaughan, & Y. Dauphin (Eds.), *35th conference on advances in neural information processing systems* (Vol. arXiv:2104.04646).
- Jacques, B. G., Tiganj, Z., Sarkar, A., Howard, M., & Sederberg, P. (2022). A deep convolutional neural network that is invariant to time rescaling. In *International conference on machine learning* (pp. 9729–9738).
- Jansson, Y., & Lindeberg, T. (2021). Scale-invariant scale-channel networks: Deep networks that generalise to previously unseen scales. *arXiv preprint arXiv:2106.06418*.
- Jeong, H., Taylor, A., Floeder, J. R., Lohmann, M., Mihalas, S., Wu, B., ... Namboodiri, V. M. K. (2022). Mesolimbic dopamine release conveys causal associations. *Science*, eabq6740.
- Jin, D. Z., Fujii, N., & Graybiel, A. M. (2009). Neural representation of time in cortico-basal ganglia circuits. *Proceedings of the National Academy of Sciences*, 106(45), 19156–19161.
- Jones, M. N., & Mewhort, D. J. K. (2007). Representing word meaning and order information composite holographic lexicon. *Psychological Review*, 114, 1–32.
- Ke, N. R., Goyal, A., Bilaniuk, O., Binas, J., Mozer, M. C., Pal, C., & Bengio, Y. (2018). Sparse attentive backtracking: Temporal credit assignment through reminding. In *Advances in neural information processing systems* (pp. 7640–7651).
- Khona, M., & Fiete, I. R. (2021). Attractor and integrator networks in the brain. *arXiv preprint arXiv:2112.03978*.
- Kim, H., Homann, J., Tank, D. W., & Berry, M. J. (2019). A long timescale stimulus history effect in the primary visual cortex. *bioRxiv*, 585539.
- King, J., & Dehaene, S. (2014). Characterizing the dynamics of mental representations: the temporal generalization method. *Trends in cognitive sciences*, 18(4), 203–210.
- Koay, S. A., Charles, A. S., Thiberge, S. Y., Brody, C. D., & Tank, D. W. (2022). Sequential and efficient neural-population coding of complex task information. *Neuron*, 110(2), 328–349.
- Komura, Y., Tamura, R., Uwano, T., Nishijo, H., Kaga, K., & Ono, T. (2001). Retrospective and prospective coding for predicted reward in the sensory thalamus. *Nature*, 412(6846), 546–9. doi: 10.1038/35087595
- Kraus, B. J., Robinson, R. J., 2nd, White, J. A., Eichenbaum, H., & Hasselmo, M. E. (2013). Hippocampal “time cells”: time versus path integration. *Neuron*, 78(6), 1090–101. doi: 10.1016/j.neuron.2013.04.015
- Kurth-Nelson, Z., & Redish, A. D. (2009). Temporal-difference reinforcement learning with distributed representations. *PLoS One*, 4(10), e7362.
- Lerner, Y., Honey, C. J., Katkov, M., & Hasson, U. (2014). Temporal scaling of neural responses to compressed and dilated natural speech. *Journal of neurophysiology*, 111(12), 2433–2444.
- Lindeberg, T., & Fagerström, D. (1996). Scale-space with casual time direction. In *European conference on computer vision* (pp. 229–240).
- Lubenov, E. V., & Siapas, A. G. (2009). Hippocampal theta oscillations are travelling waves. *Nature*, 459(7246), 534–9.
- Ludvig, E. A., Sutton, R. S., & Kehoe, E. J. (2008). Stimulus representation and the timing of reward-prediction errors in models of dopamine system. *Neural Computation*, 20, 3034–3054.
- MacDonald, C. J., Lepage, K. Q., Eden, U. T., & Eichenbaum, H. (2011). Hippocampal “time cells” bridge the gap in memory for discontinuous events. *Neuron*, 71(4), 737–749.
- Mainen, Z. F., & Kepecs, A. (2009). Neural representation of behavioral outcomes in the orbitofrontal cortex. *Curr Opin Neurobiol*, 19(1), 84–91. doi: 10.1016/j.conb.2009.03.010
- Masset, P., Malik, A. N., Kim, H. R., Bech Vilaseca, P., & Uchida, N. (2022). A distributional code for learning across timescales in dopamine-based reinforcement learning. In *Society for neuroscience abstracts* (Vol. 234.27).
- Mello, G. B., Soares, S., & Paton, J. J. (2015). A scalable population code for time in the striatum. *Current Biology*, 25(9),

- 1113–1122.
- Miri, A., Bhasin, B. J., Aksay, E. R., Tank, D. W., & Goldman, M. S. (2022). Oculomotor plant and neural dynamics suggest gaze control requires integration on distributed timescales. *The Journal of Physiology*, 600(16), 3837–3863.
- Momennejad, I., & Howard, M. W. (2018). Predicting the future with multi-scale successor representations. *bioRxiv*, 449470.
- Morcos, A. S., & Harvey, C. D. (2016). History-dependent variability in population dynamics during evidence accumulation in cortex. *Nature Neuroscience*, 19(12), 1672–1681.
- Murdock, B. B. (1982). A theory for the storage and retrieval of item and associative information. *Psychological Review*, 89, 609–626.
- Murray, J. D., Bernacchia, A., Roy, N. A., Constantinidis, C., Romo, R., & Wang, X.-J. (2017, dec). Stable population coding for working memory coexists with heterogeneous neural dynamics in prefrontal cortex. *Proceedings of the National Academy of Sciences*, 114(2), 394–399. Retrieved from <https://doi.org/10.1073/pnas.1619449114> doi: 10.1073/pnas.1619449114
- Nakahira, Y., Liu, Q., Sejnowski, T. J., & Doyle, J. C. (2021). Diversity-enabled sweet spots in layered architectures and speed–accuracy trade-offs in sensorimotor control. *Proceedings of the National Academy of Sciences*, 118(22).
- Namoodiri, V. M. K., Otis, J. M., van Heeswijk, K., Voets, E. S., Alghorazi, R. A., Rodriguez-Romaguera, J., ... Stuber, G. D. (2019). Single-cell activity tracking reveals that orbitofrontal neurons acquire and maintain a long-term memory to guide behavioral adaptation. *Nature neuroscience*, 22(7), 1110–1121.
- Ning, W., Bladon, J. H., & Hasselmo, M. E. (2022). Complementary representations of time in the prefrontal cortex and hippocampus. *Hippocampus*.
- Ogata, K. (1970). Describing-function analysis of nonlinear control systems. *Modern Control Engineering*, 2, 645–676.
- Palmer, S. E., Marre, O., Berry, M. J., 2nd, & Bialek, W. (2015). Predictive information in a sensory population. *Proceedings of the National Academy of Sciences USA*, 112(22), 6908–13. doi: 10.1073/pnas.1506855112
- Pastalkova, E., Itskov, V., Amarasingham, A., & Buzsáki, G. (2008). Internally generated cell assembly sequences in the rat hippocampus. *Science*, 321(5894), 1322–7.
- Patel, J., Fujisawa, S., Berényi, A., Royer, S., & Buzsáki, G. (2012). Traveling theta waves along the entire septotemporal axis of the hippocampus. *Neuron*, 75(3), 410–7. doi: 10.1016/j.neuron.2012.07.015
- Phillips, J. W., Schulmann, A., Hara, E., Winnubst, J., Liu, C., Valakh, V., ... others (2019). A repeated molecular architecture across thalamic pathways. *Nature neuroscience*, 22(11), 1925–1935.
- Piantadosi, S. T. (2016). A rational analysis of the approximate number system. *Psychonomic Bulletin & Review*, 1–10.
- Plate, T. A. (1995). Holographic reduced representations. *IEEE Transactions on Neural Networks*, 6(3), 623–641.
- Rainer, G., Rao, S. C., & Miller, E. K. (1999). Prospective coding for objects in primate prefrontal cortex. *Journal of Neuroscience*, 19(13), 5493–5505.
- Rao, R. P., & Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1), 79–87. doi: 10.1038/4580
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory*. New York: Appleton-Century-Crofts.
- Roitman, J. D., & Shadlen, M. N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *Journal of neuroscience*, 22(21), 9475–9489.
- Rossi-Pool, R., Zizumbo, J., Alvarez, M., Vergara, J., Zainos, A., & Romo, R. (2019). Temporal signals underlying a cognitive process in the dorsal premotor cortex. *Proceedings of the National Academy of Sciences*, 116, 7523–7532.
- Roy, D. S., Zhang, Y., Halassa, M. M., & Feng, G. (2022). Thalamic subnetworks as units of function. *Nature Neuroscience*, 25(2), 140–153.
- Russ, B. E., Koyano, K. W., Day-Cooney, J., Perwez, N., & Leopold, D. A. (2022). Temporal continuity shapes visual responses of macaque face patch neurons. *bioRxiv*.
- Sarel, A., Finkelstein, A., Las, L., & Ulanovsky, N. (2017). Vectorial representation of spatial goals in the hippocampus of bats. *Science*, 355(6321), 176–180.
- Schoenbaum, G., Chiba, A. A., & Gallagher, M. (1998). Orbitofrontal cortex and basolateral amygdala encode expected outcomes during learning. *Nature neuroscience*, 1(2), 155–159.
- Schoenbaum, G., & Roesch, M. (2005). Orbitofrontal cortex, associative learning, and expectancies. *Neuron*, 47(5), 633–6. doi: 10.1016/j.neuron.2005.07.018
- Schonhaut, D. R., Aghajan, Z. M., Kahana, M. J., & Fried, I. (2022). A neural code for spatiotemporal context. *bioRxiv*.
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, 275, 1593–1599.
- Scott, B. B., Constantinople, C. M., Akrami, A., Hanks, T. D., Brody, C. D., & Tank, D. W. (2017). Fronto-parietal cortical circuits encode accumulated evidence with a diversity of timescales. *Neuron*, 95(2), 385–398.
- Shahbaba, B., Li, L., Agostinelli, F., Saraf, M., Cooper, K. W., Haghverdian, D., ... Fortin, N. J. (2022). Hippocampal ensembles represent sequential relationships among an extended sequence of nonspatial events. *Nature communications*, 13(1), 1–17.
- Shankar, K. H., & Howard, M. W. (2010). Timing using temporal context. *Brain Research*, 1365, 3–17.
- Shankar, K. H., & Howard, M. W. (2013). Optimally fuzzy temporal memory. *Journal of Machine Learning Research*, 14, 3753–3780.
- Shankar, K. H., Singh, I., & Howard, M. W. (2016). Neural mechanism to simulate a scale-invariant future. *Neural Computation*, 28, 2594–2627.
- Shikano, Y., Ikegaya, Y., & Sasaki, T. (2021). Minute-encoding neurons in hippocampal-striatal circuits. *Current Biology*, 31(7), 1438–1449.
- Solstad, T., Boccara, C. N., Kropff, E., Moser, M. B., & Moser, E. I. (2008). Representation of geometric borders in the entorhinal cortex. *Science*, 322(5909), 1865–8.
- Spitmaan, M., Seo, H., Lee, D., & Soltani, A. (2020). Multiple timescales of neural dynamics and integration of task-relevant signals across cortex. *Proceedings of the National Academy of Sciences*, 117(36), 22522–22531.
- Steinberg, J., & Sompolinsky, H. (2022). Associative memory of structured knowledge. *bioRxiv*.
- Sutton, R. S. (1988). Learning to predict by the methods of temporal differences. *Machine learning*, 3(1), 9–44.
- Sutton, R. S., & Barto, A. G. (1981). Toward a modern theory of adaptive networks: Expectation and prediction. *Psychological Review*, 88, 135–171.

- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT Press.
- Tano, P., Dayan, P., & Pouget, A. (2020). A local temporal difference code for distributional reinforcement learning. *Advances in Neural Information Processing Systems*, *33*, 13662–13673.
- Taxidis, J., Pnevmatikakis, E. A., Dorian, C. C., Mylavarapu, A. L., Arora, J. S., Samadian, K. D., ... Golshani, P. (2020). Differential emergence and stability of sensory and temporal representations in context-specific hippocampal sequences. *Neuron*, *108*(5), 984–998.e9.
- Terada, S., Sakurai, Y., Nakahara, H., & Fujisawa, S. (2017). Temporal and rate coding for discrete event sequences in the hippocampus. *Neuron*, *94*, 1-15.
- Tiganj, Z., Cromer, J. A., Roy, J. E., Miller, E. K., & Howard, M. W. (2018). Compressed timeline of recent experience in monkey IPFC. *Journal of Cognitive Neuroscience*, *30*, 935-950.
- Tiganj, Z., Gershman, S. J., Sederberg, P. B., & Howard, M. W. (2019). Estimating scale-invariant future in continuous time. *Neural Computation*, *31*(4), 681–709.
- Tiganj, Z., Kim, J., Jung, M. W., & Howard, M. W. (2017). Sequential firing codes for time in rodent mPFC. *Cerebral Cortex*, *27*, 5663-5671.
- Tiganj, Z., Singh, I., Esfahani, Z. G., & Howard, M. W. (2022). Scanning a compressed ordered representation of the future. *Journal of Experimental Psychology: General*, *151*(3082–3096).
- Tsao, A., Sugar, J., Lu, L., Wang, C., Knierim, J. J., Moser, M.-B., & Moser, E. I. (2018). Integrating time from experience in the lateral entorhinal cortex. *Nature*, *561*, 57-62.
- Tsao, A., Yousefzadeh, S. A., Meck, W. H., Moser, M.-B., & Moser, E. I. (2022). The neural bases for timing of durations. *Nature Reviews Neuroscience*, 1–20.
- van der Meer, M. A. A., & Redish, A. D. (2011). Theta phase precession in rat ventral striatum links place and reward information. *Journal of Neuroscience*, *31*(8), 2843-54. doi: 10.1523/JNEUROSCI.4869-10.2011
- Waelti, P., Dickinson, A., & Schultz, W. (2001). Dopamine responses comply with basic assumptions of formal learning theory. *Nature*, *412*(6842), 43-8.
- Wagner, M. J., Kim, T. H., Savall, J., Schnitzer, M. J., & Luo, L. (2017). Cerebellar granule cells encode the expectation of reward. *Nature*, *544*(7648), 96–100.
- Wagner, M. J., & Luo, L. (2020). Neocortex–cerebellum circuits for cognitive processing. *Trends in neurosciences*, *43*(1), 42–54.
- Watkins, C. J., & Dayan, P. (1992). Q-learning. *Machine learning*, *8*(3-4), 279–292.
- Wei, W., Mohebi, A., & Berke, J. (2021). Striatal dopamine pulses follow a temporal discounting spectrum. *bioRxiv*.
- Wei, X.-X., & Stocker, A. A. (2012). Efficient coding provides a direct link between prior and likelihood in perceptual bayesian inference. *Advances in neural information processing systems*, *25*.
- Wilson, M. A., & McNaughton, B. L. (1993). Dynamics of the hippocampal ensemble code for space. *Science*, *261*, 1055-8.
- Young, J. J., & Shapiro, M. L. (2011). The orbitofrontal cortex and response selection. *Ann N Y Acad Sci*, *1239*, 25-32. doi: 10.1111/j.1749-6632.2011.06279.x
- Yu, Q., Bi, Z., Jiang, S., Yan, B., Chen, H., Wang, Y., ... Zhang, J. (2022). Visual cortex encodes timing information in humans and mice. *Neuron*, *110*, 4194-4211.
- Zhang, H., & Jacobs, J. (2015). Traveling theta waves in the human hippocampus. *Journal of Neuroscience*, *35*(36), 12477-87. doi: 10.1523/JNEUROSCI.5102-14.2015