

Encoding Time in Feedforward Trajectories of a Recurrent Neural Network Model

N. F. Hardy

nhardy01@gmail.com

*Neuroscience Interdepartmental Program and Department of Neurobiology,
University of California Los Angeles, Los Angeles, CA 90095, U.S.A.*

Dean V. Buonomano

dbuono@ucla.edu

*Neuroscience Interdepartmental Program and Departments of Neurology and
Psychology, University of California Los Angeles, Los Angeles, CA 90095, U.S.A.*

Brain activity evolves through time, creating trajectories of activity that underlie sensorimotor processing, behavior, and learning and memory. Therefore, understanding the temporal nature of neural dynamics is essential to understanding brain function and behavior. In vivo studies have demonstrated that sequential transient activation of neurons can encode time. However, it remains unclear whether these patterns emerge from feedforward network architectures or from recurrent networks and, furthermore, what role network structure plays in timing. We address these issues using a recurrent neural network (RNN) model with distinct populations of excitatory and inhibitory units. Consistent with experimental data, a single RNN could autonomously produce multiple functionally feedforward trajectories, thus potentially encoding multiple timed motor patterns lasting up to several seconds. Importantly, the model accounted for Weber's law, a hallmark of timing behavior. Analysis of network connectivity revealed that efficiency—a measure of network interconnectedness—decreased as the number of stored trajectories increased. Additionally, the balance of excitation (E) and inhibition (I) shifted toward excitation during each unit's activation time, generating the prediction that observed sequential activity relies on dynamic control of the E/I balance. Our results establish for the first time that the same RNN can generate multiple functionally feedforward patterns of activity as a result of dynamic shifts in the E/I balance imposed by the connectome of the RNN. We conclude that recurrent network architectures account for sequential neural activity, as well as for a fundamental signature of timing behavior: Weber's law.

1 Introduction

The ability to accurately tell time and generate appropriately timed motor responses is essential to most forms of sensory and motor processing. However, the neural processes used to encode time remain unknown (Mauk & Buonomano, 2004; Buhusi & Meck, 2005; Ivry & Schlerf, 2008; Merchant, Harrington, & Meck, 2013). While the brain tells time across many scales, ranging from microseconds to days, it is on the scale of tens of milliseconds to a few seconds that timing is most relevant to sensory-motor processing and behavior. Several neural mechanisms have been proposed to account for temporal processing in this range (for reviews, see Hardy & Buonomano, 2016; Hass & Durstewitz, 2016), including pacemaker and counter internal clocks (Treisman, 1963; Gibbon, Church, & Meck, 1984), ramping firing rates (Durstewitz, 2003; Simen, Balci, deSouza, Cohen, & Holmes, 2011), the duration of firing rate increases (Gavornik, Shuler, Loewenstein, Bear, & Shouval, 2009; Namboodiri, Huertas, Monk, Shouval, & Hussain Shuler, 2015), models that rely on the inherent stochasticity of sensory signals and neural responses (Ahrens & Sahani, 2008, 2011), and finally “population clocks,” in which timing is encoded in the evolving patterns of activity within recurrent circuits (Buonomano & Mauk, 1994; Mauk & Donegan, 1997; Medina & Mauk, 1999; Buonomano & Laje, 2010).

The theory that time is encoded in the dynamics of large populations of neurons has received experimental support in several brain regions, including the cortex (Crowe, Averbek, & Chafee 2010; Merchant, Zarco, Pérez, Prado, & Bartolo, 2011; Harvey, Coen, & Tank, 2012; Kim, Ghim, Lee, & Jung, 2013; Crowe, Zarco, Bartolo, & Merchant, 2014; Bakhurin et al., 2017), basal ganglia (Jin, Fujii, & Graybiel, 2009; Gouvea et al., 2015; Mello, Soares, & Paton, 2015), hippocampus (Pastalkova, Itskov, Amarasingham, & Buzsaki, 2008; MacDonald, Lepage, Eden, & Eichenbaum, 2011; Kraus, Robinson, White, Eichenbaum, & Hasselmo, 2013; Modi, Dhawale, & Bhalla, 2014), and area HVC in songbirds (Hahnloser, Kozhevnikov, & Fee, 2002; Long, Jin, & Fee, 2010). Some of these studies report relatively simple, apparently feedforward, sequential patterns of activity in brain regions containing significant recurrent connectivity. A fundamental question is whether these patterns of activity are generated by truly feedforward circuits or, rather, by recurrent circuits generating “functionally feedforward” patterns of activity (Banerjee, Series, & Pouget, 2008; Goldman, 2009). Here we define “functionally feedforward trajectories” as those generated by recurrent neural networks and characterized by sequential patterns of activation (“moving bumps”) in which any given unit fires only once during a pattern.

Synfire chains are perhaps the simplest network-based model that could account for the reports of functionally feedforward patterns of activity. Typically synfire chain models consist of many pools of neurons connected in a feedforward manner such that activation of one pool results in the sequential activation of each downstream pool (Abeles, 1991; Diesmann, Gewaltig,

& Aertsen, 1999). However, cortical circuits, where functionally feedforward activity is often observed, are characterized by recurrent connections and local inhibition, features that standard synfire chain models generally lack (Harvey et al., 2012). Moreover, the capacity of these synfire networks is limited because any given neuron generally participates in only one pattern (Herrmann, Hertz, & Prügel-Bennett, 1995). To address these issues, we use a model of recurrent neural networks (RNNs) to examine how they might produce functionally feedforward patterns of activity that encode time.

Previous studies of timing using RNNs have not sought to simulate experimentally observed patterns of neural activity and have used RNNs that do not follow Dale's law. We expand on previous work (Laje & Buonomano, 2013; Rajan, Harvey Christopher, & Tank, 2016) by training RNNs that follow Dale's law to emulate experimentally observed activity patterns. In addition, unlike standard RNN models, the networks in this study have only positive value firing rates. The networks are trained using the innate-training learning rule to autonomously produce stable activity for up to 5 seconds, two orders of magnitude greater than the time constant of the units (Laje & Buonomano, 2013; Rajan et al., 2016). Our results demonstrate that RNNs can robustly encode time by generating functionally feedforward patterns of activity. Importantly, these networks account for a characteristic of motor timing known as Weber's law (Gibbon, 1977; Gibbon, Malapani, Dale, & Gallistel, 1997), and can encode multiple feedforward patterns. Analysis of trained networks revealed changes in the balance of excitation and inhibition that account for the production of this feedforward activity, thus generating an experimentally testable prediction.

2 Results

2.1 Recurrent Neural Networks Produce Functionally Feedforward Trajectories. We first examined if recurrent neural networks can generate the sequential patterns of activity observed in the cortex and hippocampus (Pastalkova et al., 2008; Crowe et al., 2010, 2014; MacDonald et al., 2011). Typically in these areas, any neuron participating in a sequence is active for periods of hundreds of milliseconds to a few seconds. We used a modified version of standard firing-rate RNNs in which units are sparsely and randomly connected to one another (Sompolinsky, Crisanti, & Sommers, 1988). Specifically, to more closely mimic neural physiology, we incorporated separate populations of excitatory and inhibitory units (see Figure 1A). Furthermore, the firing rate of each unit was bounded between 0 and 1 (see section 4).

We used a supervised learning rule to adjust the recurrent weights and train the network to produce a functionally feedforward trajectory in response to a brief (50 ms) input. Specifically, the networks were trained to produce a 5 sec target sequence of feedforward activity such that each unit in the network was transiently activated without being driven by external

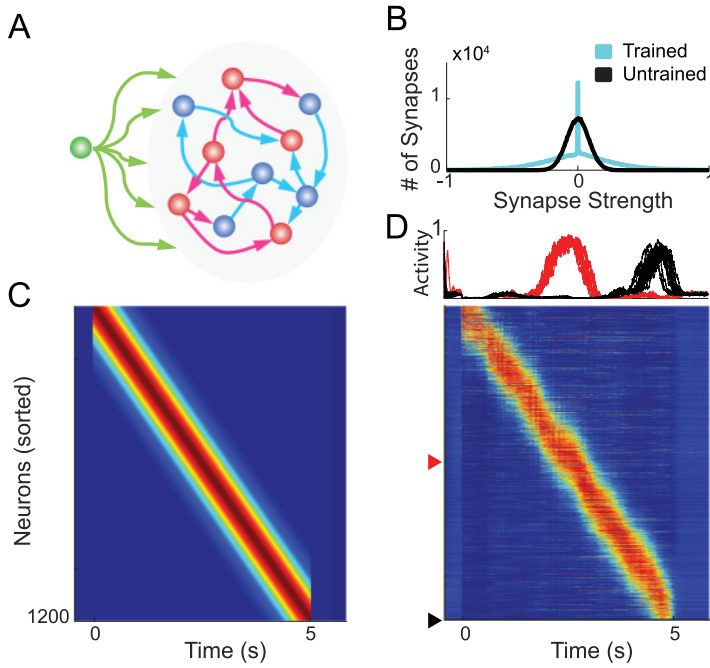


Figure 1: Generation of feedforward trajectories within a RNN. (A) Schematic of network architecture. The networks were composed of 600 excitatory (blue) and 600 inhibitory (red) firing rate units ($N^{Rec} = 1200$), with sparse recurrent connections. Neurons at the beginning of a sequence received input (green) during a 50 ms window to trigger the trajectory. (B) Connection weights were initialized with a gaussian distribution. After training to produce a single feedforward sequence, a large number of weights were pruned to zero, and some weights became stronger, resulting in a long-tailed distribution. (C) Example five sec feedforward sequence target. (D) After training, the RNN can produce a 5 sec feedforward trajectory. Top: Two units trained to activate in the middle and end of the trajectory, highlighted below. Each trace represents 1 of 15 trials. Bottom: Example network activity from a single trial.

input (see Figure 1C). This activity pattern can be thought of as a moving bump of neural activity. After training, the network was able to reproduce the 5 sec neural trajectory in response to the brief input (see Figure 1D). Importantly, after the end of the trajectory, the RNN returned to an inactive rest state. Thus, in contrast to RNNs in high-gain regimes, these networks were silent at rest. Training dramatically altered the distribution of synaptic weights in the network: the weight of many synapses converged to 0 (in part as a consequence of the boundaries imposed by Dale's law) while others were strengthened, resulting in long tails (see Figure 1B). These

long tails of the synaptic weight distribution are in line with experimentally observed distributions of synaptic weights (Song, Sjöström, Reigl, Nelson, & Chklovskii, 2005) and observations in previous models of neural dynamics in RNNs (Laje & Buonomano, 2013; Rajan et al., 2016).

2.2 RNNs Can Encode Multiple Sequences. Many motor behaviors such as playing the piano or writing require the use of the same muscle groups activated in distinct temporal patterns. If the motor cortex is to drive these motor patterns, it must produce distinct well-timed trajectories of neural activity using the same sets of neurons activated in different orders. Traditional models of sequential neural activity (e.g., standard feedforward synfire chains) do not account for this because each unit generally participates in only a single sequence.

To examine the capacity of recurrent networks to encode multiple functionally feedforward trajectories, we trained RNNs to learn patterns in which all units participated in each trajectory. RNNs were trained to learn 1, 3, 5, 10, or 20 distinct sequences. Each sequence lasted 1 sec and was triggered by a distinct input. As shown in Figure 2A, an RNN can generate multiple distinct patterns in response to distinct inputs. Importantly, each pattern recruits all the units in the network. To quantify the network capacity, we calculated the correlation between the evoked activity on each test trial and the corresponding target. We used the average correlation across targets as a measure of performance. Trained networks could reliably produce 10 1 sec sequences with relatively little decrement in performance. However, when RNNs were trained on 20 patterns, they showed a large decrease in performance (one-way ANOVA, $F_{4,45} = 193$, $p < 10^{-27}$, $n = 10$ networks; see Figure 2B) and increased failure rate (number of trials in which the input did not evoke a pattern or generated a partial sequence; one-way ANOVA, $F_{4,45} = 419$, $p < 10^{-35}$, $n = 10$ networks).

2.3 Functionally Feedforward Trajectories Account for Weber's Law. A defining feature of behavioral timing is that there is an approximately linear relationship between the standard deviation and mean of a timed response (Gibbon, 1977; Gibbon et al., 1997), referred to as the scalar property or Weber's law. The ability to account for Weber's law is often taken as a benchmark for models of timing and does not generally emerge spontaneously in many models (Ahrens & Sahani, 2008; Hass & Herrmann, 2012; Hass & Durstewitz, 2014). To examine whether the RNNs studied here obey Weber's law, we measured the temporal variability of each unit within a single feedforward pattern at different levels of noise. We fit the activity of each unit on every test trial with a gaussian function and calculated the standard deviation and mean of the peak time of each unit's fit across trials (see section 4). We used Weber's generalized law to fit the standard deviation as a function of time and refer to the slope of this linear fit as the Weber coefficient (Ivry & Hazeltine, 1995; Merchant et al., 2008). Note that

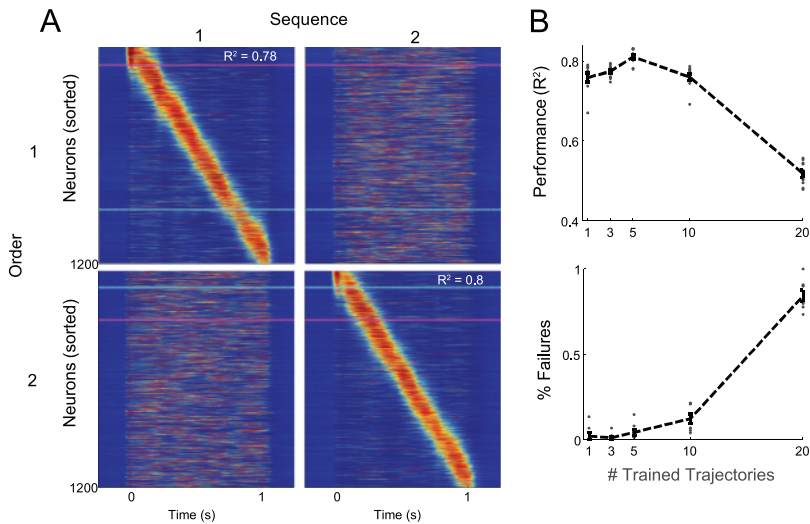


Figure 2: A single RNN can encode many different feedforward trajectories. (A) The units of the same trained RNN in response to two different inputs. Each column shows the spatiotemporal pattern of activity triggered by a single input. Each row shows the activity sorted according to feedforward trajectory 1 (top row) or 2 (bottom row). Blue and pink dashed lines highlight the same two units in all panels. (B) Performance (top) and failure rate (bottom) for 10 networks trained to produce up to 20 trajectories. Each dot represents the average across 15 trials per target for a single network. Each network can reliably produce up to 10 feedforward sequences before performance decreases. Error bars show the mean \pm SEM.

Weber's generalized law allows for a positive intercept. In each of 10 trained networks, we found that the standard deviation of a unit's peak firing time across trials increased linearly with its mean activation time (see Figure 3A). This property was highly robust: while the Weber coefficient increased with the amount of noise injected into the network, the scalar property was preserved even at large noise amplitudes (see Figure 3B). Thus, Weber's law was not limited to any specific noise parameter choice.

2.4 Sequences Are Generated by Dynamic Shifts in the Balance of Excitation and Inhibition. Experimental studies have conclusively demonstrated that functionally feedforward patterns of activity occur *in vivo*, but these studies have not been able to explore the neural and network mechanisms underlying these patterns (Pastalkova et al., 2008; Harvey et al., 2012; MacDonald, Carrow, Place, & Eichenbaum, 2013). Computational models allow us to address this question and generate experimental predictions.

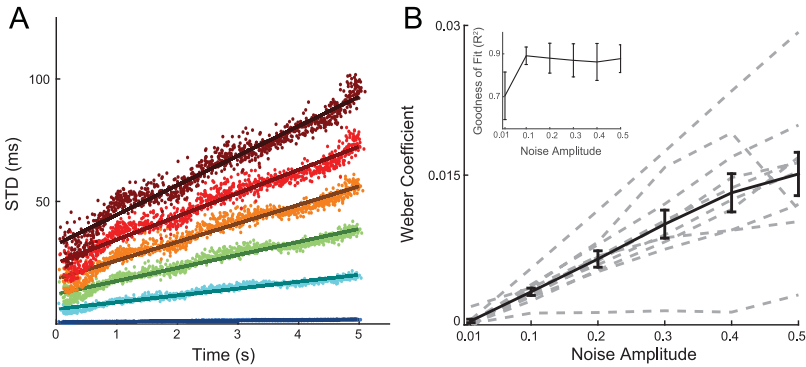


Figure 3: Timing generated by the RNN obeys Weber's law. (A) Temporal variability increases linearly with elapsed time. Each dot shows the standard deviation of the activation time of a unit of an example network plotted against its mean. The activation time of each unit on each trial was determined from the center of a gaussian fit of that unit's activity. Each color represents one of six amplitudes of injected noise. Lines show the linear fit. (B) The Weber law is robust to noise. The Weber coefficient (slope of the linear fit shown in panel A) is shown across six noise amplitudes for 10 trained networks. Each dotted line represents a single network, and the solid black line represents the mean with error bars showing \pm SEM. Inset: Mean \pm SEM of the goodness of fit (R^2) for the linear fits.

To determine how RNNs generate sequential activity, we first examined the balance of excitation and inhibition in the units during the trained patterns. This analysis parallels experimental studies that have examined the relative balance of excitatory and inhibitory currents (Shu, Hasenstaub, & McCormick, 2003; Froemke, Merzenich, & Schreiner, 2007; Heiss, Katz, Ganmor, & Lampl, 2008). We examined the balance of excitation and inhibition by separately summing the total excitatory and inhibitory input onto each unit at all time points. Figure 4A shows the relationship between the E/I balance and firing rate in a single example unit. Interestingly, when cells were not active, they still received significant excitation and inhibition, and these inputs were usually approximately balanced ($E/I \cong 1$), similar to a recent *in vitro* study of temporal processing (Goel & Buonomano, 2016). This balance shifted toward excitation ($E/I > 1$) primarily during a unit's target activity period. By plotting the E/I ratio of all the neurons in the network during a trajectory, it is possible to visualize the progressive shift in the E/I balance during a feedforward sequence (see Figure 4B). These observations are consistent with experimental studies during up-state activity showing that excitatory and inhibitory currents are balanced and that changes in activity consist of subtle shifts toward excitation (Shu et al., 2003; Haider, Duque, Hasenstaub, & McCormick, 2006; Sun et al., 2010), as well as recent work

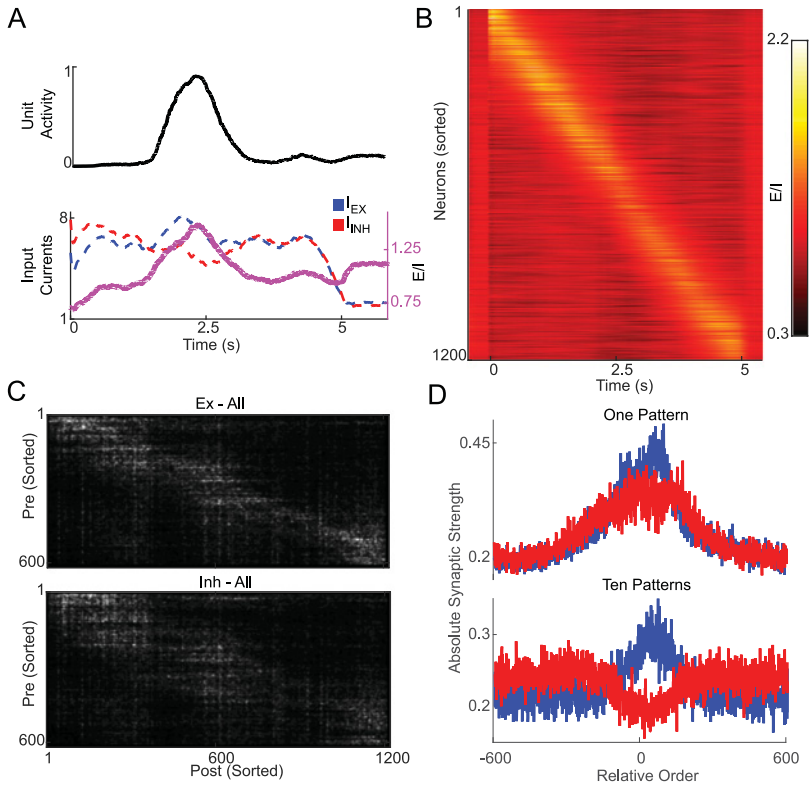


Figure 4: Feedforward trajectories are generated by dynamic shifts in the E/I balance. (A) The excitation/inhibition (E/I) ratio peaks around a unit's activation time. Top: The activity of an example unit, which peaked in the middle of a 5 sec trajectory. Bottom: Inhibitory (red) and excitatory (blue) inputs are high but usually balanced throughout the trajectory, but the balance shifts toward excitation when the unit is active. E/I ratio shown in pink. (B) A heat-map of the ratio of the same network showing a shifting peak of the ratio along the trajectory (during one trial). (C) Excitatory (top) and inhibitory (bottom) weight matrices of a network sorted according to the activation order. Note the peak in synaptic strength following the activation order along the diagonal. Weights are smoothed for visualization. (D) Averaging along the feedforward trajectory reveals a peak in excitatory weights pointing forward, bounded by inhibition. The absolute value of the weight matrix was centered along the feedforward sequence and averaged according to excitatory and inhibitory units. Top: Weights of a network trained for one pattern. Bottom: The same initial network, but trained for 10 feedforward patterns.

showing timing-specific shifts in E/I balance in vitro (Goel & Buonomano, 2016).

To examine the origins of the dynamic shift in the E/I balance, we analyzed the synaptic connectivity matrix. When sorted according to activation order, connections from both inhibitory and excitatory units displayed a pattern of peak connection strength along the sequence of activation (see Figure 4C). To examine the structure of the connectome, we shifted the sorted weight matrix to align the window of activity of each postsynaptic cell within the trajectory. Taking the mean across all cells of this shifted weight matrix revealed a peak of excitation pointing forward along the trajectory (i.e., the excitatory weights are asymmetrically shifted to the right), bounded by peaks of inhibition (see Figure 4D, upper panel). Despite the rightward shift of the peak—and in contrast to feedforward networks—the excitatory units are clearly connected in the “forward” and “backward” directions. This anatomical feature allows for the local mutual excitation necessary to keep units active for durations of up to 1 sec. This Mexican hat connectivity pattern has been observed in other studies of sequence generation (Itskov, Curto, Pastalkova, & Buzsáki, 2011; Rajan et al., 2016) and accounts for the moving bump of activation in feedforward RNNs. Interestingly, for a single pattern, the shift of the E/I balance toward excitation was primarily driven by an increase in excitation (see Figure 4D, upper panel). However, when the same analysis is performed for networks that learned 10 patterns, the E/I shift driving activity forward was generated by both an increase in excitation and a decrease in inhibition (see Figure 4D, lower panel). Taken together, these results predict that recurrent networks in the cortex generate functionally feedforward sequences of activity using asymmetries in the connectivity patterns between neurons and dynamic shifts in the E/I balance.

2.5 Connectivity of RNNs Reflects the Number of Encoded Trajectories. There is increasing emphasis on characterizing the microcircuit structure, or the connectome, of biological neural circuits. To further examine the relationship between the microcircuit structure in our model and potentially generate experimental predictions, we calculated network efficiency—a standard measure from graph theory that captures the interconnectedness of the units in a graph (Boccaletti, Latora, Moreno, Chavez, & Hwang, 2006). Specifically, it measures the minimal weighted path length between units, such that a larger efficiency value corresponds to a shorter path (see section 4). Because we were interested in the relationship between structure and function, we compared efficiency measures of coactive and noncoactive units. Furthermore, since the activity in any unit is the result of the interaction between excitatory and inhibitory inputs, we separately calculated the net excitatory and inhibitory connection strengths between pairs of units, generating topological representations of recurrent excitation

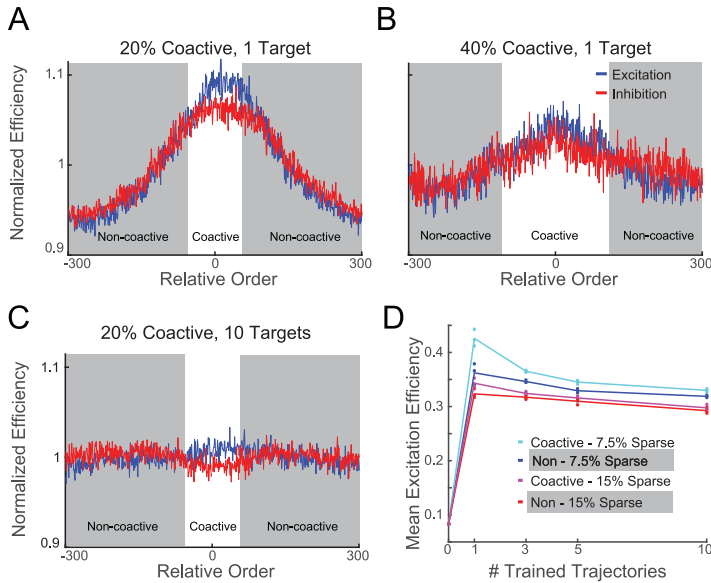


Figure 5: Multiple feedforward patterns are embedded via uniform path lengths between units. (A) Average disynaptic efficiency between coactive (unshaded) and noncoactive (shaded) unit pairs in an RNN trained for a single feedforward target. Units that are coactive during a trajectory have a higher mean efficiency compared to noncoactive pairs. Efficiency values were normalized to the mean to aid visualization. (B) Same as panel A but showing the same initial network trained such that more units are coactive at a given time (lower temporal sparsity). The efficiency across the network is more uniform. (C.) Same as in panel A, now trained for 10 targets. In panels A to C, postsynaptic efficiency values were aligned according to the activation order of the presynaptic unit within the same target order and averaged across units. (D) Average efficiency of connections between coactive and noncoactive units according to the number of trained trajectories. Efficiency increases sharply from the naive weights. As more trajectories are encoded, efficiency becomes uniform across the network, that is, the difference between coactive and noncoactive units decreases.

and inhibition. We then calculated the weighted efficiency index of these connections in trained untrained networks.

As expected, when we averaged along a single trained sequence of feedforward activity, we observed that units that coactivate (i.e., active at neighboring points in time) had a higher-than-average efficiency value (see Figure 5A), similar to the observed Mexican hat architecture in Figure 4D. However, when networks were trained to sustain more coactive units (see Figure 5B), or were trained to produce a larger number of targets (see Figure 5C), the connection efficiency between coactive and noncoactive

units approached the network mean. Finding the average disynaptic efficiency between coactive and noncoactive units revealed that efficiency sharply increased when networks were trained for a single target, with coactive efficiency exceeding noncoactive (Figure 5D). As more trajectories were encoded, this difference decreased, indicating that efficiency became uniform with respect to unit pairs' active relationship. Moreover, networks with more coactive units (40% active) were initially more uniform than those with fewer coactive units (i.e., the efficiency between coactive and noncoactive units was more similar), consistent with the notion that higher local efficiency may be necessary to maintain temporally sparser trajectories in order to support more local positive feedback between coactive units.

3 Discussion

Functionally feedforward patterns of activity have been observed in a wide range of different brain areas (Pastalkova et al., 2008; MacDonald et al., 2011; Kraus et al., 2013; Mello et al., 2015; Bakhurin et al., 2017). These patterns have been proposed to underlie a number of different behaviors, including memory, planning, and motor timing. Here we have focused primarily on the potential role of such patterns in timing, specifically in tasks that animals must learn to generate timed motor patterns or anticipate when external events will occur on the scale of hundreds of milliseconds to seconds. Our results show that even though many of these experimentally reported patterns of sequential activation are apparently accounted for by feedforward architectures, recurrent neural networks are more consistent with the data. Furthermore, recurrent architectures are computationally more powerful in that they can store many different trajectories in which each unit participates in each trajectory. We propose that networks with recurrent excitation underlie the functionally feedforward trajectories observed in cortical areas.

A number of models have proposed mechanisms for generating functionally feedforward patterns within recurrent networks (Buonomano, 2005; Liu & Buonomano, 2009; Fiete, Senn, Wang, & Hahnloser, 2010; Itskov et al., 2011; Rajan et al., 2016). These studies have used both spiking and firing rate models and relied on a number of different mechanisms, but they have not explicitly addressed a standard benchmark for behavioral timing: Weber's law. One recent model developed by Rajan et al. (2016) also trained RNNs using an RLS-based learning rule, and our results complement their findings: specifically, by tuning the recurrent weights of an initially randomly connected network, it is possible to robustly encode multiple functionally feedforward patterns of activity. That study, however, focused primarily on sequence generation and encoding memory-dependent motor behaviors; it did not encode time per se as the network was driven in part by time-varying inputs (i.e., external information about the time from trial onset was present).

3.1 Weber's Law. Here we show that an RNN can encode time and account for Weber's law, more specifically Weber's generalized law, which states that the standard deviation of a timer increases linearly with elapsed time (Ivry & Hazeltine, 1995; Merchant et al., 2008; Laje, Cheng, & Buonomano, 2011). The origins of Weber's law in timing models is a long-standing and vexing problem because, according to the simplest model in which an accumulator integrates the pulses of a noisy oscillator, the standard deviation of the latency of a neuron, or of a motor output, should increase as a function of the square root of total time (Hass & Herrmann, 2012; Hass & Durstewitz, 2014, 2016). In contrast, the current model naturally captures Weber's law (at least within the parameter regimes used here), even at high noise amplitudes. Major issues remain, such as the properties underlying Weber's law in recurrent networks and why the brain settles for the observed linear relationship (Hass & Herrmann, 2012). We hypothesize that these properties may be related: recurrency may inherently amplify internal noise, producing long-lasting temporal correlations (Hass & Herrmann, 2012), and evolutionarily speaking, the trade-off was adaptive because it increased computational capacity.

The current model also establishes that RNNs can robustly store multiple patterns, in which each neuron participates in every pattern. This feature is consistent with experimental findings demonstrating that the same neuron can participate in multiple patterns of network activity, firing within different windows in each (Pastalkova et al., 2008; MacDonald et al., 2013). Thus, the experimental data and the current model are consistent with stimulus-specific timing in which time codes are generated in relation to each stimulus or task condition as opposed to an absolute time code. The capacity of the RNNs described here appears to be fairly large. But as demonstrated in a previous study, the true capacity of RNNs is likely to be strongly dependent on model assumptions, most notably noise levels (Laje & Buonomano, 2013).

Within this population clock framework, the same RNN does not function as a single clock, but rather implements many event-specific timers. That is, the network does not encode absolute time but elapsed time from stimulus onset, and there is an entirely different time code for each stimulus. This computational strategy ensures that the activity vector at any given instant not only encodes elapsed time but also provides a dynamic memory of the current stimulus.

3.2 RNN Connectome. In order for sequences to propagate in a defined trajectory through a network, activity must generate imbalances that simultaneously push the activity forward and prevent it from deviating from the proper activation order (Ben-Yishai, Bar-Or, & Sompolinsky, 1995; Fiete et al., 2010). Here we find that training an RNN composed of distinct excitatory and inhibitory populations produces synaptic connectivity resembling an asymmetric Mexican hat architecture, with excitation propagating

and maintaining network activity and inhibition bounding this activity to prevent off-target activation. Importantly, the recurrency of the network enables multiple Mexican hats to be embedded in a single connectivity matrix, allowing multiple functional feedforward patterns to be produced by a single network.

An important characteristic of the connectome of a network is how efficiently individual units exchange information. Surprisingly, we found that the weighted efficiency of a feedforward RNN was negatively correlated with the number of sequences stored and that this change was largely driven by reduced efficiency between coactive units. Indeed, multitrajectory networks exhibited uniform path lengths between units, regardless of their relative activation order. This flattening of the efficiency within a network is likely necessary to allow units that are highly separated in one sequence to also co-activate in another. Thus, a prediction that emerges from this study is that learning may induce an overall decrease in the efficiency of cortical circuits as the networks embed more uniform connection structures, making individually learned patterns difficult to distinguish using connectomics.

Reports of sequential patterns of activity in multiple brain areas appear to be superficially consistent with feedforward synfire-like architectures. However, recurrent networks are likely responsible for generating the experimentally observed patterns for two reasons. First, although the patterns of activity comprise sequential activation of neurons, the duration of activity over which a neuron fires (in the range of hundreds of millisecond to a few seconds) likely relies on local positive feedback maintained by recurrent connections, and, second, purely feedforward network architectures are unlikely to account for the ability of networks to generate multiple trajectories in which any given neuron can participate in many different patterns.

4 Materials and Methods

4.1 Network Structure and Dynamics. The network dynamics were governed by the standard firing rate equations (Abeles, 1982; Sompolinsky et al., 1988; Jaeger & Haas, 2004):

$$\tau \frac{dx_i}{dt} = -x_i + \sum_{j=1}^{N^{Rec}} W_{ij}^{Rec} r_j + \sum_{k=1}^{N^{In}} W_{ik}^{In} y_k + I_i^{Noise}, \quad (4.1)$$

where x_i represents the state of unit i . The sparse, $N^{Rec} \times N^{Rec}$ matrix W^{Rec} describes the recurrent connectivity, with nonzero values initially drawn from a gaussian distribution of $\mu = 0$ and $\sigma = g/\sqrt{p_c * N^{Rec}}$, where $g = 1.6$ is the synaptic scaling constant and $p_c = 0.3$ is the probability a given unit

will connect to another unit in the network (autapses were eliminated). The firing rate, r_i , of unit i is given by the logistic function,

$$r_i = \frac{1}{1 + e^{-a \cdot x_i + b}}, \quad (4.2)$$

where $a = 2$ and $b = 4$ correspond to the gain and threshold of the units, respectively. Compared to the traditional tanh function, this provides a more biologically plausible model in which activity is low at rest (i.e., without input), and rates are bounded between 0 and 1.

After initialization, the efferent synapses from a randomly selected half of the recurrent units were set to be positive and the other half set to be negative to create inhibitory and excitatory populations. The $N^{Rec} \times N^{In}$ matrix W^{In} describes the input connections from input units y to the recurrent units and is set to stimulate only those units active at the start of a functionally feedforward sequence with weight equal to each unit's target activity at $t = 0$. The activity of the input units was set to 0 except during the 50 ms input window at the beginning of a trial when it was set to 3. $N^{Rec} = 1200$ is the network size, and $\tau = 25$ ms is the time constant of the units. The random noise current (I^{Noise}) was drawn from a gaussian distribution of $\mu = 0$ and $\sigma = 0.5$ (noise amplitude) unless otherwise indicated. A single unit of the RNN contacted all other units and was tonically active, providing a bias to each unit but containing no temporal information because its activity was set to 1 at all times.

4.2 Training the RNNs. The RNNs were trained to generate target patterns of sequential activity designed to mimic the functionally feedforward activity observed in neural circuits during temporal tasks (Hahnloser et al., 2002; Pastalkova et al., 2008; MacDonald et al., 2011; Harvey et al., 2012). These targets were generated by setting each unit to activate briefly in sequence so the entire population tiled the interval defined by t_{max} (see Figure 1C). The activation order was generated randomly for each target pattern and constrained so that the order of inhibitory and excitatory units was interleaved. The pattern of activation for each unit was set by a gaussian with a μ equal to the unit's activation time. The temporal sparsity of the pattern, defined by $\frac{N^{Active}}{N^{Rec}}$, where N^{Active} represents the number of units in the target pattern that are active at any given point in time, was set to approximately 20% by making the σ of the gaussian target function 7.5% of t_{max} (Rajan et al., 2016).

Training was performed using the innate training learning rule, which tunes the recurrent weights based on errors generated by each unit (Laje & Buonomano, 2013), similar to the learning rule used in Rajan et al. (2016). The error was determined by taking the difference between the activity of the unit r_i and its target activity at time t and used to update its weights

using the recursive least squares (RLS) algorithm (see also Haykin, 2002; Jaeger & Haas, 2004; Sussillo & Abbott, 2009; Mante, Sussillo, Shenoy, & Newsome, 2013; Carnevale, de Lafuente, Romo, Barak, & Parga, 2015; Rajan et al., 2016). As described previously (Laje & Buonomano, 2013) the weights onto a given unit were updated proportional to its error, the activity of its presynaptic units, and the inverse cross-correlation matrix of the network activity. To maintain Dale's law, efferent weights from all units were bounded so that they could not cross zero (i.e., negative weights were prevented from becoming positive and vice versa). The weights were bounded to a maximum value of $g + 5/\sqrt{N^{Rec}}$ of the appropriate sign to prevent overfitting. Training was conducted at a noise amplitude of 0.5, and all recurrent units were trained.

4.3 Weber Analysis. To determine if the timing of the network obeyed Weber's law, we tested each network 15 times at six different I^{Noise} amplitudes. For each trial, we fit the activity of each unit with a gaussian curve and used the center of that curve as a measure of the unit's activation time. Because each unit activates only once, time can be measured directly from the activity state of the network (Abeles, 1982; Long et al., 2010). Only units whose gaussian fit had an $R^2 > 0.9$ were used for the Weber analysis. For a given noise level and network, we calculated the standard deviation (std_i) and mean (t_i) of these activation times for each unit i and found a linear fit of these values. We then used the slope of the linear fit as the Weber coefficient, std/t (excluding units outside a 95% confidence interval of the linear fit; see Figure 1).

4.4 Performance. Network performance was measured by a performance index, calculated as the correlation (R^2) between network activity on a given trial and the corresponding target pattern. The overall performance of a network was calculated as the network's mean performance index across all trials for all trained patterns. Particularly at high noise levels and numbers of trained targets, networks sometimes failed to complete a feedforward sequence; thus, we also used the percentage of these failures to quantify capacity (Figure 2B).

4.5 Network Efficiency. In graph theory, network efficiency measures the shortest path between two nodes of a network and can be thought of as a measure of interconnectedness of the units (Boccaletti et al., 2006). Efficiency was calculated by determining the minimum weighted disynaptic excitatory and inhibitory path length between pairs of excitatory units in the network, with weights normalized to the maximum weight. Disynaptic connection strengths were calculated by taking matrix products $W_{ExEx}W_{ExEx}$ for the excitatory path and $W_{ExInh}W_{InhEx}$ for the inhibitory path, creating two $N^{Ex} \times N^{Ex}$ matrices. The path length between two units was determined by finding the series of edges that connected the units with the

smallest summed inverse weight. For example, if unit A is connected to unit B with a strength of 0.2, the path length will be 5; if unit A is also connected to unit C with a strength 0.5 and unit C connects to B with a strength of 0.5, then the length from A to B is 4 ($2 + 2$). Thus, the minimum weighted path between A and B would be through unit C. This path length was calculated for all possible pairs in the disynaptic matrices, inverted, and normalized by the total number of possible connections ($N^{Rec} * (N^{Rec} - 1)$) to generate an efficiency value. Therefore, a network with 100% maximal connections would have an efficiency value of 1.

Acknowledgments

We thank Vishwa Goudar for his helpful comments on the manuscript. This research was supported by NIH grants MH60163 and T32 NS058280.

References

- Abeles, M. (1982). *Local cortical circuits: An electrophysiological study*. Berlin: Springer.
- Abeles, M. (1991). *Corticonics*. Cambridge: Cambridge University Press.
- Ahrens, M. B., & Sahani, M. (2008). Inferring elapsed time from stochastic neural processes. In J. C. Platt, D. Koller, Y. Singer, & S. T. Roweis (Eds.), *Advances in neural information processing systems, 20* (pp. 1–8). Red Hook, NY: Curran.
- Ahrens, M. B., & Sahani, M. (2011). Observers exploit stochastic models of sensory change to help judge the passage of time. *Current Biology, 21*, 200–206.
- 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. *J. Neurosci, 37*, 854–870.
- Banerjee, A., Series, P., & Pouget, A. (2008). Dynamical constraints on using precise spike timing to compute in recurrent cortical networks. *Neural Comput., 20*, 974–993.
- Ben-Yishai, R., Bar-Or, R. L., & Sompolinsky, H. (1995). Theory of orientation tuning in visual cortex. *Proceedings of the National Academy of Sciences, 92*, 3844–3848.
- Boccaletti, S., Latora, V., Moreno, Y., Chavez, M., & Hwang, D. U. (2006). Complex networks: Structure and dynamics. *Physics Reports, 424*, 175–308.
- Buhusi, C. V., & Meck, W. H. (2005). What makes us tick? Functional and neural mechanisms of interval timing. *Nature Reviews Neuroscience, 6*, 755–765.
- Buonomano, D. V. (2005). A learning rule for the emergence of stable dynamics and timing in recurrent networks. *J. Neurophysiol, 94*, 2275–2283.
- Buonomano, D. V., & Laje, R. (2010). Population clocks: Motor timing with neural dynamics. *Trends in Cognitive Sciences, 14*, 520–527.
- Buonomano, D. V., & Mauk, M. D. (1994). Neural network model of the cerebellum: Temporal discrimination and the timing of motor responses. *Neural Comput., 6*, 38–55.
- Carnevale, F., de Lafuente, V., Romo, R., Barak, O., & Parga, N. (2015). Dynamic control of response criterion in premotor cortex during perceptual detection under temporal uncertainty. *Neuron, 86*, 1067–1077.

- Crowe, D. A., Averbeck, B. B., & Chafee, M. V. (2010). Rapid sequences of population activity patterns dynamically encode task-critical spatial information in parietal cortex. *Journal of Neuroscience*, *30*, 11640–11653.
- Crowe, D. A., Zarco, W., Bartolo, R., & Merchant, H. (2014). Dynamic representation of the temporal and sequential structure of rhythmic movements in the primate medial premotor cortex. *Journal of Neuroscience*, *34*, 11972–11983.
- Diesmann, M., Gewaltig, M. O., & Aertsen, A. (1999). Stable propagation of synchronous spiking in cortical neural networks. *Nature*, *402*, 529–533.
- Durstewitz, D. (2003). Self-organizing neural integrator predicts interval times through climbing activity. *Journal of Neuroscience*, *23*, 5342–5353.
- Fiete, I. R., Senn, W., Wang, C. Z. H., & Hahnloser, R. H. R. (2010). Spike-time-dependent plasticity and heterosynaptic competition organize networks to produce long scale-free sequences of neural activity. *Neuron*, *65*, 563–576.
- Froemke, R. C., Merzenich, M. M., & Schreiner, C. E. (2007). A synaptic memory trace for cortical receptive field plasticity. *Nature*, *450*, 425–429.
- Gavornik, J. P., Shuler, M. G. H., Loewenstein, Y., Bear, M. F., & Shouval, H. Z. (2009). Learning reward timing in cortex through reward dependent expression of synaptic plasticity. *Proceedings of the National Academy of Sciences*, *106*, 6826–6831.
- Gibbon, J. (1977). Scalar expectancy theory and Weber's law in animal timing. *Psychological Review*, *84*, 279–325.
- Gibbon, J., Church, R. M., & Meck, W. H. (1984). Scalar timing in memory. *Ann. NY Acad. Sci.*, *423*, 52–77.
- Gibbon, J., Malapani, C., Dale, C. L., & Gallistel, C. R. (1997). Toward a neurobiology of temporal cognition: Advances and challenges. *Curr. Opinion Neurobiol.*, *7*, 170–184.
- Goel, A., & Buonomano, D. V. (2016). Temporal interval learning in cortical cultures is encoded in intrinsic network dynamics. *Neuron*, *91*, 320–327.
- Goldman, M. S. (2009). Memory without feedback in a neural network. *Neuron*, *61*, 621–634.
- Gouvea, T. S., Monteiro, T., Motiwala, A., Soares, S., Machens, C., & Paton, J. J. (2015). Striatal dynamics explain duration judgments. *Elife*, *4*.
- Hahnloser, R. H. R., Kozhevnikov, A. A., & Fee, M. S. (2002). An ultra-sparse code underlies the generation of neural sequence in a songbird. *Nature*, *419*, 65–70.
- Haider, B., Duque, A., Hasenstaub, A. R., & McCormick, D. A. (2006). Neocortical network activity in vivo is generated through a dynamic balance of excitation and inhibition. *J. Neurosci.*, *26*, 4535–4545.
- Hardy, N. F., & Buonomano, D. V. (2016). Neurocomputational models of interval and pattern timing. *Current Opinion in Behavioral Sciences*, *8*, 250–257.
- Harvey, C. D., Coen, P., & Tank, D. W. (2012). Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature*, *484*, 62–68.
- Hass, J., & Durstewitz, D. (2014). Neurocomputational models of time perception. In H. Merchant & V. de Lafuente (Eds.), *Neurobiology of interval timing* (pp. 49–73). New York: Springer.
- Hass, J., & Durstewitz, D. (2016). Time at the center, or time at the side? Assessing current models of time perception. *Current Opinion in Behavioral Sciences*, *8*, 238–244.

- Hass, J., & Herrmann, J. M. (2012). The neural representation of time: An information-theoretic perspective. *Neural Computation, 24*, 1519–1552.
- Haykin, S. (2002). *Adaptive filter theory*. Upper Saddle River, NJ: Prentice Hall.
- Heiss, J. E., Katz, Y., Ganmor, E., & Lampl, I. (2008). Shift in the balance between excitation and inhibition during sensory adaptation of S1 neurons. *J. Neurosci., 28*, 13320–13330.
- Herrmann, M., Hertz, J. A., & Prügel-Bennett, A. (1995). Analysis of synfire chains. *Network: Computation in Neural Systems, 6*, 403–414.
- Itskov, V., Curto, C., Pastalkova, E., & Buzsáki, G. (2011). Cell assembly sequences arising from spike threshold adaptation keep track of time in the hippocampus. *Journal of Neuroscience, 31*, 2828–2834.
- Ivry, R. B., & Hazeltine, R. E. (1995). Perception and production of temporal intervals across a range of durations: Evidence for a common timing mechanism. *Journal of Experimental Psychology—Human Perception and Performance, 21*, 3–18.
- Ivry, R. B., & Schlerf, J. E. (2008). Dedicated and intrinsic models of time perception. *Trends in Cognitive Sciences, 12*, 273–280.
- Jaeger, H., & Haas, H. (2004). Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication. *Science, 304*, 78–80.
- Jin, D. Z., Fujii, N., & Graybiel, A. M. (2009). Neural representation of time in cortico-basal ganglia circuits. *Proc. Natl. Acad. Sci. USA, 106*, 19156–19161.
- Kim, J., Ghim, J.-W., Lee, J. H., & Jung, M. W. (2013). Neural correlates of interval timing in rodent prefrontal cortex. *Journal of Neuroscience, 33*, 13834–13847.
- Kraus, B. J., Robinson, R. J., White, J. A., Eichenbaum, H., & Hasselmo, M. E. (2013). Hippocampal “time cells”: Time versus path integration. *Neuron, 78*, 1090–1101.
- Laje, R., & Buonomano, D. V. (2013). Robust timing and motor patterns by taming chaos in recurrent neural networks. *Nat. Neurosci., 16*, 925–933.
- Laje, R., Cheng, K., & Buonomano, D. V. (2011). Learning of temporal motor patterns: An analysis of continuous versus reset timing. *Frontiers in Integrative Neuroscience, 5*, 61.
- Liu, J. K., & Buonomano, D. V. (2009). Embedding multiple trajectories in simulated recurrent neural networks in a self-organizing manner. *Journal of Neuroscience, 29*, 13172–13181.
- Long, M. A., Jin, D. Z., & Fee, M. S. (2010). Support for a synaptic chain model of neuronal sequence generation. *Nature, 468*, 394–399.
- MacDonald, C. J., Carrow, S., Place, R., & Eichenbaum, H. (2013). Distinct hippocampal time cell sequences represent odor memories in immobilized rats. *J. Neurosci., 33*, 14607–14616.
- 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*, 737–749.
- Mante, V., Sussillo, D., Shenoy, K. V., & Newsome, W. T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature, 503*, 78–84.
- Mauk, M. D., & Buonomano, D. V. (2004). The neural basis of temporal processing. *Annual Review of Neuroscience, 27*, 307–340.
- Mauk, M. D., & Donegan, N. H. (1997). A model of Pavlovian eyelid conditioning based on the synaptic organization of the cerebellum. *Learn. Mem., 3*, 130–158.

- Medina, J. F., & Mauk, M. D. (1999). Simulations of cerebellar motor learning: Computational analysis of plasticity at the mossy fiber to deep nucleus synapse. *Journal of Neuroscience*, *19*, 7140–7151.
- Mello, G. B. M., Soares, S., & Paton, J. J. (2015). A scalable population code for time in the striatum. *Curr. Biol.*, *9*, 1113–1122.
- Merchant, H., Zarco, W., & Prado, L. (2008). Do we have a common mechanism for measuring time in the hundreds of millisecond range? Evidence from multiple-interval timing tasks. *J. Neurophysiol.*, *99*, 939–949.
- Merchant, H., Harrington, D. L., & Meck, W. H. (2013). Neural basis of the perception and estimation of time. *Annual Review of Neuroscience*, *36*, 313–336.
- Merchant, H., Zarco, W., Pérez, O., Prado, L., & Bartolo, R. (2011). Measuring time with different neural chronometers during a synchronization-continuation task. *Proceedings of the National Academy of Sciences*, *108*, 19784–19789.
- Modi, M. N., Dhawale, A. K., & Bhalla, U. S. (2014). CA1 cell activity sequences emerge after reorganization of network correlation structure during associative learning. *eLife*, *3*, e01982.
- Namboodiri, V. M. K., Huertas, M. A., Monk, K. J., Shouval, H. Z., Hussain, & Shuler, M. G. (2015). Visually cued action timing in the primary visual cortex. *Neuron*, *86*, 319–330.
- Pastalkova, E., Itskov, V., Amarasingham, A., & Buzsaki, G. (2008). Internally generated cell assembly sequences in the rat hippocampus. *Science*, *321*, 1322–1327.
- Rajan, K., Harvey Christopher, D., & Tank, D. W. (2016). Recurrent network models of sequence generation and memory. *Neuron*, *90*, 128–142.
- Shu, Y., Hasenstaub, A., & McCormick, D. A. (2003). Turning on and off recurrent balanced cortical activity. *Nature*, *423*, 288–293.
- Simen, P., Balci, F., deSouza, L., Cohen, J. D., & Holmes, P. (2011). A model of interval timing by neural integration. *Journal of Neuroscience*, *31*, 9238–9253.
- Sompolinsky, H., Crisanti, A., & Sommers, H. J. (1988). Chaos in random neural networks. *Physical Review Letters*, *61*, 259–262.
- Song, S., Sjöström, P. J., Reigl, M., Nelson, & Chklovskii, D. B. (2005). Highly non-random features of synaptic connectivity in local cortical circuits. *PLoS Biology*, *3*, 508–518.
- Sun, Y. J., Wu, G. K., Liu, B.-H., Li, P., Zhou, M., Xiao, Z., . . . Zhang, L. I. (2010). Fine-tuning of pre-balanced excitation and inhibition during auditory cortical development. *Nature*, *465*, 927–931.
- Sussillo, D., & Abbott, L. F. (2009). Generating coherent patterns of activity from chaotic neural networks. *Neuron*, *63*, 544–557.
- Treisman, M. (1963). Temporal discrimination and the indifference interval: Implications for a model of the “internal clock.” *Psychol. Monogr.*, *77*, 1–31.