J Neurophysiol 94: 2275–2283, 2005;
doi:10.1152/jn.01250.2004.

A Learning Rule for the Emergence of Stable Dynamics and Timing in Recurrent Networks

Dean V. Buonomano

Departments of Neurobiology and Psychology and Brain Research Institute, University of California, Los Angeles, Los Angeles, California

Submitted 6 December 2004; accepted in final form 18 May 2005

Buonomano, Dean V. A learning rule for the emergence of stable dynamics and timing in recurrent networks. *J Neurophysiol* 94: 2275–2283, 2005; doi:10.1152/jn.01250.2004. Neural dynamics within recurrent cortical networks is an important component of neural processing. However, the learning rules that allow networks composed of hundreds or thousands of recurrently connected neurons to develop stable dynamical states are poorly understood. Here I use a neural network model to examine the emergence of stable dynamical states within recurrent networks. I describe a learning rule that can account both for the development of stable dynamics and guide networks to states that have been observed experimentally, specifically, states that instantiate a sparse code for time. Across trials, each neuron fires during a specific time window; by connecting the neurons to a hypothetical set of output units, it is possible to generate arbitrary spatial-temporal output patterns. Intertrial jitter of the spike time of a given neuron increases as a direct function of the delay at which it fires. These results establish a learning rule by which cortical networks can potentially process temporal information in a self-organizing manner, in the absence of specialized timing mechanisms.

INTRODUCTION

Local dynamics within cortical networks play a fundamental role in neural computations (Koch and Fuster 1989; Ringach et al. 1997; Somers et al. 1995). Dynamics in the form of complex spatial-temporal patterns of neuronal firing have been observed in vitro (Beggs and Plenz 2004; Buonomano 2003; Ikegaya et al. 2004; Jimbo et al. 1999) and in vivo and been shown to contain information about sensory stimuli (Gawne et al. 1996; Laurent et al. 1996) and motor behavior (Hahnloser et al. 2002; Weesberg et al. 2000). While it has been established that networks exhibit complex spatial-temporal dynamics, the synaptic learning rules that allow recurrent networks to develop functional and stable dynamics are not known. Indeed, conventional coincident-based learning rules are unstable (Miller and MacKay 1994; Turrigiano and Nelson 2004) and can lead to runaway excitation in recurrent networks.

A potential computational function of spatial-temporal patterns of activity is temporal processing and motor control. Decoding temporal information, or generating timed responses on the order of tens to hundreds of milliseconds, is a fundamental component of many sensory and motor tasks (Mauk and Buonomano 2004). Indeed the precise sequential generation of motor responses is a virtually ubiquitous component of behavior. One of the most studied forms of complex sensory-motor processing is the birdsong system (Bottjer and Arnold 1997; Doupe and Kuhl 1999). Song generation relies on precisely timed sequential generation of motor patterns over both the time scale of individual syllable features and sequences of syllables (Fee et al. 2004). While relatively little is known about the neural mechanisms underling the generation of precisely timed motor sequences, it has recently been shown that there is a sparse code for time in the premotor area HVc, which may control song production (Hahnloser et al. 2002). Dynamically changing patterns of activity have also been proposed to code for time in the cerebellum and underlie certain motor patterns (Medina et al. 2000). Sparse long-lasting responses have also been observed in vitro. As shown in Fig. 1 in cortical organotypic cultures, a single stimulus can elicit single spikes at latencies of a few hundred milliseconds (Buonomano 2003). These slices contain thousands of recurrently connected neurons and initially exhibit weak synaptic connections (Echevarría and Albus 2000; Muller et al. 1993). In this study, I examined how dynamics may emerge in this general class of networks.

It has been shown that a modified form of synaptic scaling can guide networks to a stable dynamical state. Interestingly, the dynamical states that emerge from this learning rule allow the network to generate a sparse temporal code for time. Qualitatively the architecture of the model and dynamical states observed are consistent with experimental data from cortical organotypic slices (Fig. 1) (Buonomano 2003).

METHODS

All simulations were performed with NEURON (Hines and Carnevale 1997). Each neuron was simulated as a single compartment integrate-and-fire unit, with dynamic synapses. The ratio of excitatory (Ex) to inhibitory (Inh) neurons and connectivity probability was based on experimental data (Beaulieu et al. 1992). Specifically, 80% of the units were excitatory and 20% inhibitory. In our default simulations, there were 320 Ex and 80 Inh units (each Ex unit received 20 excitatory and 5 inhibitory synapses, and each Inh unit received 5 excitatory synapses). Connectivity was uniformly random.

Integrate-and-fire units

The resting membrane potential of all units was −60 mV. Thresholds were set from a normal distribution ($\sigma^2 = 5\%$ of mean threshold); the mean thresholds for the Ex and Inh units were −40 and −45 mV, respectively. After a spike, the voltage was reset to −60 and −65 for the Ex and Inh units, respectively. Membrane time constants were 30 ms for the Ex units and 10 ms for the Inh units. Input

Address for reprint requests and other correspondence: D. V. Buonomano, Brain Research Inst., Univ. of California, Box 951761, Los Angeles, CA 90095 (E-mail: dbuono@ucla.edu).

The costs of publication of this article were defrayed in part by the payment of page charges. The article must therefore be hereby marked "*advertisement*" in accordance with 18 U.S.C. Section 1734 solely to indicate this fact.

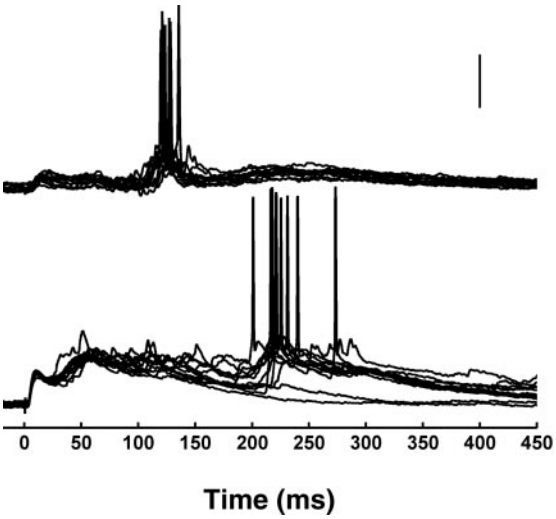

FIG. 1. Long-latency sparse responses in cortical networks in vitro. Late responses in cortical organotypic slices in response to a single electrical stimulus at *time 0*. *Top* and *bottom panels* represent 12 traces from 2 cells. Calibration bar, 20 mV. Modified from Buonomano 2003.

resistance was 300 MΩ. See supplementary information[1] for further details.

Synapses

α-Amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid (AMPA), *N*-methyl-D-aspartate (NMDA), and $GABA_A$ synaptic currents were simulated using a kinetic model (Buonomano 2000; Destexhe et al. 1994; Karmarkar and Buonomano 2002) (supplemental information). Short-term plasticity was incorporated in all the synapses based on experimental data and implemented according to Markram et al. (1998). Specifically the Ex → Ex synapses exhibited depression—U (use of synaptic efficacy parameter) = 0.5; τ_{rec} = 700 ms—Ex → Inh synapses exhibited facilitation: U = 0.2; τ_{rec} = 125 ms; τ_{fac} = 500 ms. Inh → Ex dynamics were based on basket cell synapses: U = 0.25; τ_{rec} = 700 ms; τ_{fac} = 25 ms (Gupta et al. 2000). The presence of short-term plasticity was not essential to the behavior of the networks described below. Initial synaptic strengths were chosen from a normal distribution. For the Ex → Ex, Ex → Inh, and Inh → Ex synapses, the mean initial synaptic strength was 0.5, 2, and 8 μS, respectively. To insure that plasticity did not result in an unphysiological state in which a single presynaptic neuron could fire a postsynaptic neuron, the maximal synaptic weight for the Ex → Ex and Ex → Inh connections was set at 10 and 30 μS, respectively. The NMDA current contributed to long latency responses but was not essential for the stability or dynamics of the network. The strength of the NMDA component was a fixed ratio of the AMPA strength (0.3). The strength of the inhibitory synapses was fixed throughout the simulations. Results were robust to large changes in initial parameters. Changes in an order of magnitude in the Inh → Ex did not alter the final homeostatic state. The learning rate α was set to 0.05 in the simulations presented here.

Synaptic scaling

The variable A_i measures the average activity of neuron i and is defined as

$$\tau \frac{dA_i(t)}{dt} = S_i(t) - A_i(t) \quad (1)$$

where $S_i(t)$ represents the presence (1) or absence (0) of a spike at time t. I arbitrarily used τ = 10 s and a stimulus presentation rate (intertrial interval) of 10 s (see DISCUSSION)

The synaptic scaling learning rule is generally represented as (van Rossum et al. 2000)

$$\Delta W_{pre,post} = \alpha \cdot [A_{goal} - A_{post}(t)] \cdot W_{pre,post} \quad (2)$$

$W_{pre,post}$ represents the synaptic strength between a presynaptic and postsynaptic neuron and is updated at the end of each trial. Unless stated otherwise, A_{goal} was set at 1 and 2 for the Ex and Inh units, respectively. A_{goal} was set higher for the Inh units, because inhibitory neurons have higher evoked and spontaneous firing rates (Swadlow 1989).

Information measure of timing accuracy

We quantified the information present in the spatial pattern of spikes of the Ex units by determining the information content of the current time bin (Borst and Theunissen 1999). In essence, we can think of each time bin as a stimulus, and ask what is the mutual information between the spike patterns (in time bin n) and time bin n. However, it was of interest to know not only the mutual information of the system over all time bins but in relation to each time bin, to determine if there is more information available at certain intervals. Thus I asked what is the information about time bin n in relation to all time bins ≠ n. The mutual information about time bin n was equal to

$$I(n) = H(S) + H(n) - H(S,n) \quad (3)$$

where $H(S)$ is the conventional entropy measure of all spike patterns, $H(n)$ reflects the entropy of being in time been n versus all other time bins, and thus is the same for all time bins and based on $p(1/N)$ and $p(1 - 1/N)$, where N is the total number of time bins. $H(S,n)$ is the joint entropy. Because the occurrence of time bin n and all other time bins are not equiprobable, the maximal mutual information at all time steps is equal to

$$I_{max} = -\left(\frac{1}{N}\right)\log_2\left(\frac{1}{N}\right) - \left(\frac{N-1}{N}\right)\log_2\left(\frac{N-1}{N}\right) \quad (4)$$

Realistic measures of timing accuracy

To obtain an estimate of the timing ability of the network as a whole that would be available to an output neuron that was trained to fire at a specific time, I incorporated five output units to the network. Each output unit received synapses from all the Ex units. On training trials, the weight between Ex_i and $output_j$ (where j is the target interval) was increased if i was active at time j. Testing was performed over a separate set of trials.

RESULTS

To address the emergence of stable dynamical states in recurrent networks, I first examined a previously described learning rule–synaptic scaling (*Eq.* 2; Turrigiano et al. 1998), which has been shown to guide neurons in feed-forward networks to a stable homeostatic state (Turrigiano and Nelson 2004; van Rossum et al. 2000). This experimentally derived learning rule was proposed in the context of maintaining approximately constant levels of activity over the course of days. Here I will examine if synaptic scaling can generalize to a different condition, specifically a condition in which activity is stimulus-driven over a shorter time course and the desired activity patterns are sparse (Hahnloser et al. 2002).

Unless otherwise stated, the artificial neural network was composed of 400 recurrently connected integrate-and-fire neurons, and synapses were initially weak. These general assump-

[1] The Supplementary Material for this article (two movies and a figure) is available online at http://jn.physiology.org/cgi/content/full/01250.2004/DC1.

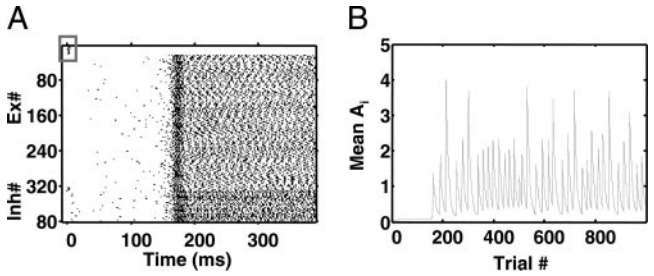

FIG. 2. Synaptic scaling leads to an unstable dynamical state. *A*: sample raster plot of the network at the end of training. Each row represents a single unit in the network. The *top* 320 rows represent excitatory (Ex) units and the *bottom* 80 rows represent inhibitory (Inh) units. A single dot represents a spike in the neuron corresponding to the unit on the *y*-axis. Small box represents subset of input neurons. *B*: network dynamics did not stabilize. Although target activity level of all Ex units was 1, the network did not achieve a homeostatic state in which neurons fired at their target rate. Rather, overall activity (Mean A_i) oscillated over 1,000 trials.

tions are meant to loosely reflect those observed in dissociated or organotypic cultures (Muller et al. 1993). The target level of activity was initially set to one spike per trial for the Ex neurons. This assumption was meant to capture the sparse activity often observed in response to a single stimulus in organotypic cultures (Fig. 1) (Buonomano 2003), as well as data from primary sensory areas indicating that neurons generally respond to a single transient stimuli with one spike (e.g., Armstrong-James et al. 1994; Kilgard and Merzenich 1998). Each trial consisted of activating a subset of neurons (input neurons), which again is comparable with extracellular stimulation in an in vitro slice, where a random subset of neurons close to the electrode will function as an input source. The network evolved under the guidance of synaptic scaling over 1,000 trials. Synaptic scaling essentially states that the conductances of all synapses onto a postsynaptic neuron are increased (multiplicatively) if the average activity level (A_i) of the cell is below a predetermined setpoint and decreased if the cell is hyperactive. Figure 2 shows the results from simulations in which the setpoint of the Ex neurons was set to 1 spike/trial (see Movie 1). In the initial trials, activation of the input neurons did not result in any suprathreshold activity in the other neurons. With training, the learning rule was effective in generating network activity. However, it did not converge to a steady state in which neurons stabilized at their target activity level. Instead, oscillatory behavior was observed. This behavior was observed in dozens of stimulations with different initial conditions and independent of the learning rate (see also Fig. 4). Because synaptic scaling is known to be stable in feed-forward networks (Turrigiano and Nelson 2004), I hypothesized the instability arises as a result of the recurrent architecture. As shown in Fig. 3*A*, using a simplified implementation of the above network, I examined the effects of recurrency. If all postsynaptic Ex neurons received only a single synapse (thus effectively implementing a feed-forward network), each neuron reached it's target level of activity and the network converged. If a minimal degree of recurrency was introduced, by assigning two or four synapses to each neuron, convergence was not observed.

An inherent limitation of synaptic scaling when incorporated into recurrent networks can be schematized in a toy model (Fig. 3*B*). Consider two reciprocally connected neurons (A and B), one of which (A) receives a synapse from an input neuron. If the synapses onto A from the input (w_1) and B (w_2) start in an initial condition such that $w_1 < w_2$, then because w_1 and w_2 will always be scaled by the same factor, w_1 will always be smaller than w_2. However, if the network is to converge to a stable state in which each neuron is active once per trial, clearly w_1 must be larger than w_2. Although this is simply a toy model, one can see a potential limitation with synaptic scaling. The ratio of all synapses onto a cell will always remain the same because Δw is a function only of postsynaptic activity. In essence, although a cell may have 1,000 synapses, it only has one degree of freedom. The network cannot converge because there is no solution that maintains the synaptic ratios onto a postsynaptic neuron. Thus the lack of convergence is independent of the parameters of the model. While the learning rates and time window over which activity is calculated can alter the

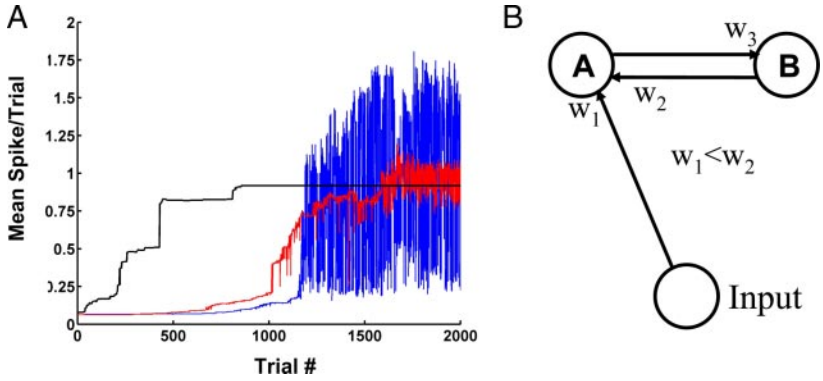

FIG. 3. Recurrency and fixed synaptic ratios contribute to the lack of convergence. *A*: average number of spikes per cell (not A_i) over 2,000 training trials, in networks in which each neurons received 1 (black), 2 (red), or 4 (blue) from other excitatory neurons. With nEx $\rightarrow$ Ex = 1, synaptic scaling converges to a state in which each neuron achieves it's target level of activity. Note that this represents a feed-forward network. Small degrees of recurrency (nEx $\rightarrow$ Ex = 2 or 4) progressively result in increasing instability. Learning rates were smaller than those in used in Fig. 2 to attempt to facilitate convergence. Note black line activity does not converge to 1, because not every neuron fired. This is because if each neuron only receives 1 synapse, some neurons will be permanently "out of the loop," e.g., neurons *Y* and *Z* could be mutually connected. In these simulations synaptic strengths were allowed to reach a strength in which a single neuron could fire another. *B*: toy model providing an example of a condition in which synaptic scaling cannot converge. If we stipulate that, in a network composed of 2 neurons, only 1 neuron receives an input, the 2nd neuron must be driven polysynaptically. If the target level of activity is 1 in both neurons, and initial weights are such that $w_1 < w_2$, homeostasis cannot be achieved with synaptic scaling. This is the case because if w_1/w_2 begins at <1, it will always be <1. By definition, synaptic scaling does not change the ratio. However, the solution clearly requires that w_1 and w_3 be above threshold and w_2 be below threshold. (Note that even if one resorts to a refractory period in cell A, the same principle holds true in a larger network chain).

magnitude and period of the oscillations, the synaptic ratios remain unchanged; thus a solution cannot be reached.

Presynaptic-dependent scaling

To allow a neuron to change the relative synaptic strength ratios during training, the synaptic scaling rule was modified to include a presynaptic component

$$\Delta W_{\text{pre,post}} = \alpha \cdot A_{\text{pre}}(t)[A_{\text{goal}} - A_{\text{post}}(t)] \cdot W_{\text{pre,post}} \quad (5)$$

Under this rule, a postsynaptic neuron will preferentially potentiate synapses from neurons that are more active (independent of *when* it was active within a trial), I will refer to *Eq.* 4 as presynaptic-dependent scaling.

Figure 4, *A* and *B*, shows the results from a simulation using the same initial parameters used in the simulations shown in Fig. 2. With training, both the neurons and global network activity converge to a stable state (Movie 2). Each Ex neuron fired once per trial, and the rate of change of synaptic weights approached zero (data not shown). Figure 2*A* also shows that there is complex spatial-temporal structure to the network dynamics. By plotting the same data sorted by latency (Fig. 4*B*; Movie 3), it is possible to observe a synfire-like pattern, in which activity propagates throughout the network. However, propagation is not based on the spatial arrangement of the neurons, because there is no topography in the network.

To compare the ability of this learning rule to drive the neurons in a recurrent network to their target level of activity with synaptic scaling, I performed simulations using different setpoints. To ensure that a true steady state was reached, 6,000 trials were performed, and the mean level of activity and stability of the last 1,000 were quantified. Figure 4*C* shows that both learning rules were effective in bringing mean activity (averaged over all Ex units over 1,000 trials) to near the setpoint. Figure 4*D* shows the SD of the mean activity level. The high SD in the synaptic scaling simulations reflects the lack of convergence. In contrast, the low SD observed using *Eq.* 4 reflects the convergence to a stable dynamical state.

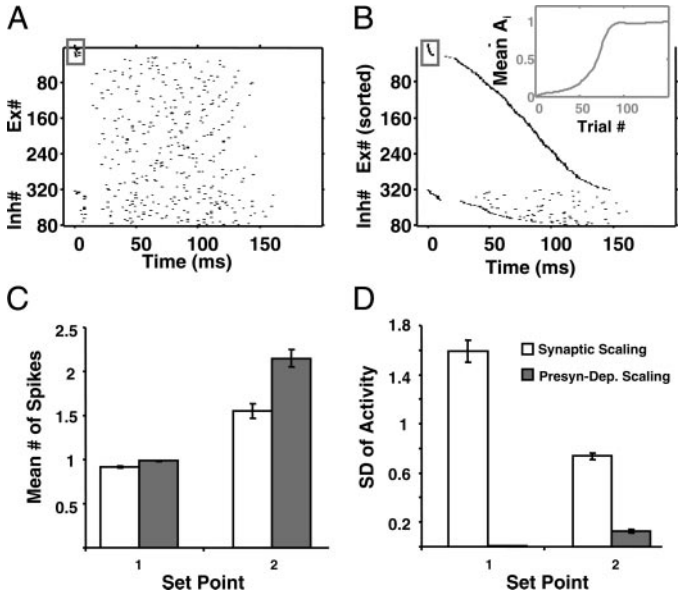

FIG. 4. Presynaptic-dependent scaling produces stable dynamics. *A*: raster plot of a network trained with presynaptic-dependent scaling. Each of the Ex units fired once and only once. *B*: same data as in *A*, with the order of Ex units determined by their latency. There is a clear spatial-temporal structure to the dynamics. *Inset*: evolution of the network through training. The *y*-axis shows average activity trace (A_i) of all neurons in the network. *C*: mean activity levels calculated over the last 1,000 of 6,000 training trials, using either synaptic scaling (empty bars) or presynaptic-dependent scaling (filled bars). Setpoint for all Ex neurons was set to either 1 or 2. Both rules produced a mean level of activity close to the setpoint. *D*: SD activity level from the same simulations shown in *C*. High SD in the synaptic scaling simulations reflects the instability of network dynamics (Fig. 1*B*), whereas low SD observed using presynaptic-dependent scaling reflects convergence of a network to a stable dynamical state (*B*).

Timing

As shown in Fig. 4*B*, presynaptic-dependent scaling produced a spatial-temporal pattern of activity, in which different neurons fired at specific time windows from stimulus onset. This pattern could potentially function as a population code for time. Figure 5*A* shows a poststimulus time histogram (PSTH) raster of all the neurons in the network after training over 25 trials. Note that across trials spike jitter is a function of latency. This is expected in a system in which spike latency variance is amplified within a trial and has been recently observed in neural networks in vitro (Buonomano 2003). The ability of the network to code for time was quantified in two ways. First, the information available in the spatial pattern of activity was calculated. Three different temporal binnings were used: 1, 4, and 8 ms. Maximal information at a given time step implies that there are spatial patterns that can be used with 100% certainty to determine the time within the accuracy of the bin used. As shown in Fig. 5*B*, at 8-ms resolution, the network can with certainty determine the time $\leq$100 ms. However, at 1-ms resolution, there are gaps. The biggest one is at $\sim$10 ms, in which no neurons are active due to fast inhibition. Thus the network cannot time with millisecond resolution. To quantify the ability of the network to tell time in a physiologically plausible manner, it was assumed that the Ex units provided input to a set of output units. Each of the output units receives inputs from all Ex units and represents an interval from 25 to 125 ms. The synaptic weights from the Ex to the output units were adjusted by training over 25 trials. Performance was subsequently tested using a distinct set of 25 trials. Figure 5*C* shows the average response of the output units during a single trial. These results show that a network without any specialized timing mechanisms can process temporal information.

Convergence time and propagation speed

The above simulations were performed with a network composed of 400 units. To determine if performance was independent of network size, I examined the performance of the presynaptic-dependent scaling learning rule with networks composed of 200, 400, 800, and 1,200 units (Fig. 6*A*). All the networks were scaled; that is, the ratio of Ex to Inh units remained constant, as did the probability of connectivity (thus the total number of synapses per neuron increased). Interestingly, the larger network converged faster. This is because, since each neuron has more synapses, the average synaptic strength for a neuron to fire is decreased. Thus fewer trials are required to allow synapses to grow sufficiently to fire a neuron. However, larger networks also tend to overshoot and thus do not converge in a monotonic fashion (for α of 0.05). Again,

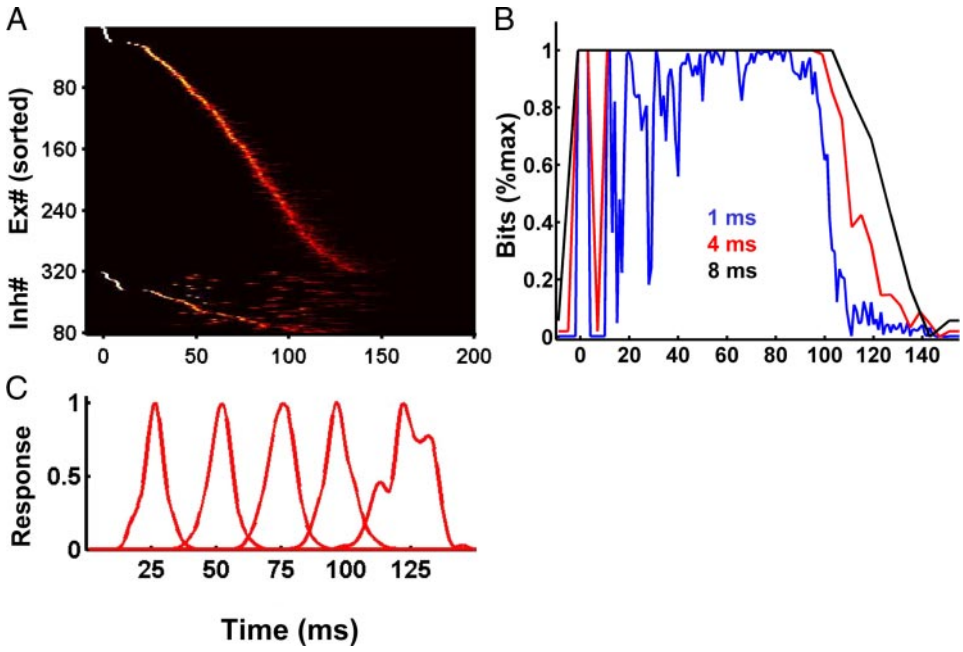

FIG. 5. Network can encode time. *A*: peristimulus time histogram (PSTH) rasters from 25 trials in the network shown in Fig. 2 after training. Color in each line represents PSTH of that cell (white corresponds to 25 spikes). Note formation of the diagonal band and increased jitter for longer latencies. *B*: information contained in the spatial pattern of Ex units about the current time bin. Three different bin sizes were examined: 1, 4, and 8 ms. A maximal value at time t indicates that, on all trials, there were patterns that would allow an observer to determine t with certainty. *C*: response of output units. All the Ex units shown in *A* were connected to 5 output units, each 1 trained to respond to 1 of 5 intervals: 25, 50, 75, 100, or 125 ms. Each trace represents response of 5 output units averaged over 25 trials. In this stimulation, output units were nonspiking, which allows for a linear measure of the structure of input.

this is because the total synaptic input grows at a faster rate because of the increased number of synapses.

We also examined the effect of changing the connection probabilities within the network. These simulations were performed in the large network with 1,200 units. The final state of three sets of simulations with three different connection probabilities are shown in Fig. 6*B*. As expected, the sparser the connectivity, the longer it takes activity to propagate throughout the network, thus extending the intervals the network could code for. However, there is a trade-off in that the intertrial variability also increases as the total duration increases (data not shown).

Lesions

To examine the effects of any single neuron on overall dynamics of the network, I performed ablation experiments (Fig. 7)—simulations in which 1% of the neurons lesioned revealed little change in the overall spatial-temporal pattern of activity. Lesioning 5% of the cells did not change the mean level of activity but did dramatically change the dynamics. Whereas most cells exhibited a longer latency, some exhibited no shift or a decrease in latency. The shift toward longer latencies is the result of decreased spatial summation, because most cells receive fewer active inputs, thus taking more time to reach threshold. This increased latency can be amplified as activity propagates throughout the network. As shown in Fig. 7*A*, some neurons can fire earlier. This can occur as a result of disinhibition: for the same reason the Ex neurons may fire later, the Inh neurons will also, allowing a few Ex neurons to fire earlier.

Coding for longer durations

In the preceding simulations, there was no spatial topography; that is, all neurons had equiprobable connection probabilities. To examine the effects of presynaptic-dependent scaling in a more realistic network, I examined the effect of topography. In these simulations, there were 1,200 units. Each Ex unit received inputs from 80 other Ex units. Connection probability was normally distributed around each neuron placed on a one-dimensional array. Figure 8 shows a raster from a single trial after training. As in the previous simulations, presynaptic-dependent scaling produced a homeostatic state in which each Ex neuron fired once during a trial. As a result of the topography, the total propagation time increased to ~350 ms. To examine the readout of a network with later responses, I trained five output units to respond to intervals of 50, 100, 150, 200, and 250 ms. Consistent with in vitro (Buonomano 2003) and psychophysical data (Karmarkar and Buonomano 2003), the across-trial jitter increased as a function of latency. These results confirm that, with or without spatial topography, presynaptic-dependent scaling produces a stable network state. Additionally, this simulation shows that the time range the network can encode is extended by using larger networks with topography. The maximal latency is extended in the topographic network because of the presence of spatial propagation: neurons distant from the input signal can only begin to fire once their neighbors have begun to fire.

DISCUSSION

The above results establish a learning rule that can lead a naïve recurrent network to achieve a stable dynamical state that implements a sparse code for time. Qualitatively, the results account for the sparse long-latency responses observed in cortical organotypic slices (Buonomano 2003).

It is well known that the control of activity in spiking recurrent networks is much more difficult then in feed-forward networks (e.g., Maass and Sontag 1999; Pearlmutter 1995). Because synaptic scaling has been shown to lead to stability in feed-forward networks (van Rossum et al. 2000) or contribute to stabilization in recurrent networks (Renart et al. 2003), our first approach was to determine if synaptic scaling could generate stable stimulus-driven dynamics. It should be noted that the conditions examined here differ from those of the

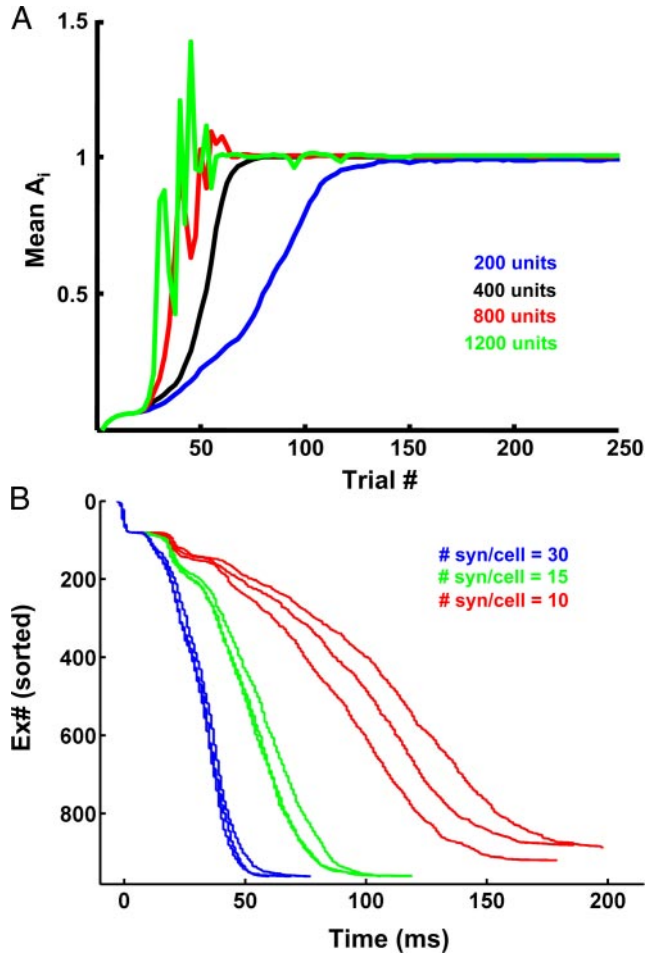

FIG. 6. Convergence as a function of network size. *A*: convergence of the presynaptic-dependent scaling rule for networks of 200, 400, 800, and 1,200 units. All networks were scaled versions of the ones presented above. Specifically, ratio of Ex to Inh units, connection probabilities, and ratio of input units were the same for all 4 simulations. *B*: speed of propagation and overall timing window is a function of connectivity density. Three simulations of 3 different networks of size 1,200 (960 Ex units), were run with 10, 15, and 30 excitatory synapses per neuron. The denser the connectivity, the quicker activity propagated throughout the network. Note that with a very sparse connectivity, not all cells could fire (see Fig. 3).

original framework of synaptic scaling. Our simulations pertain to a special case of homeostasis in which the target level of activity is very low (1 or 2 spikes per stimulus). Furthermore, the time frame I examined is on the scale of minutes and hours, whereas synaptic scaling is generally assumed to take place over days (Turrigiano and Nelson 2004). Under the specific conditions examined here, I observed that synaptic scaling by itself did not generate a stable dynamical state (Fig. 2). Note, however, that the oscillations might be considered irrelevant from the perspective of homeostatic plasticity, because averaged over a longer time period, the mean target activity level was achieved (Fig. 4). In the current simulations, instabilities arise for two reasons. First, as schematized in Fig. 3*B*, the ratio of the synaptic weights onto a neuron do not change as a result of synaptic scaling, thus preventing convergence if the solution requires a change in the synaptic ratios (Fig. 3*B*). Second, in a recurrent network, the input to a given neuron is nonstationary, a neuron *i* may find a steady-state by trial 20, and at trial 30, neuron *j,* which synapses onto *i,* may begin to fire, thus altering the behavior of *i,* which in turn can alter *j* and the network as a whole, contributing to oscillations. As shown in the toy model in Fig. 3*B*, these instabilities are not an issue of gain control. Specifically, the use of simulated annealing or an integral controller (van Rossum et al. 2000) will not change the ratio of the synaptic strengths.

Mechanisms

The learning rule described in *Eq.* 4 takes the general form of an anti-Hebbian covariance rule in which average activity, as opposed to firing rate, is used (Dayan and Abbott 2001). By averaging long time windows, the associative flavor of the traditional covariance rules that is captured by NMDA-dependent mechanisms (Brown et al. 1990), is not present. In

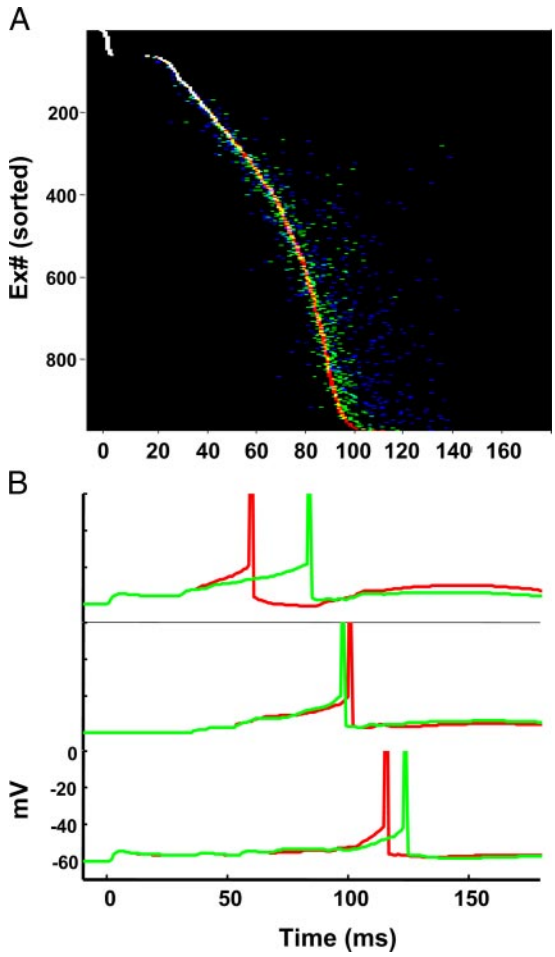

FIG. 7. Effects of cell ablation on network dynamics. *A*: raster of all the Ex units from 3 simulations from a single trial (after training) in which 0, 1, and 5% of neurons were ablated (red, green, and blue, respectively). Units were inactivated at random. White represents overlap between all 3 simulations, meaning that a neuron spiked at the same time in each simulation. With 1% of the neurons ablated, there was relatively little effect on dynamics (overlap between red and green = yellow). With 5% of cell lesioned, there was a significant change in the spatial-temporal structure of activity, generally representing increased latencies. In these stimulations, noise was removed to provide direct cross-condition comparisons. Pixel size was also increased to allow visualization of each unit, making it difficult to visually detect the inactivated units. *B*: voltage traces of 3 cells in the control and 1% lesion simulations. Traces show that different cells can exhibit different degrees of shifts in their latency. While most cells exhibit an increased latency, decreased latencies could also be observed because of disinhibition.

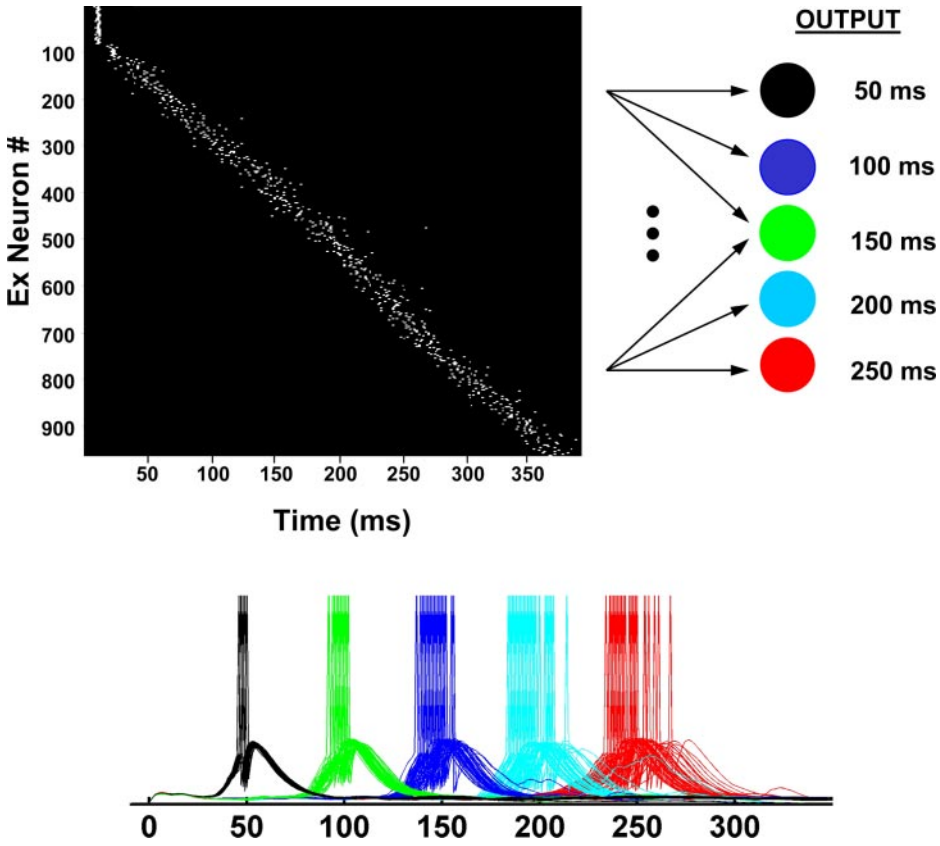

FIG. 8. Emergence of dynamics in a network with spatial topography. Dynamics in a network composed of 1,200 neurons after training with presynaptic-dependent scaling. *Top*: single trial raster after training. *Bottom*: behavior of 5 output units trained to target intervals. Training of weights onto output units was performed using a training set of 50 rasters. Output traces show response of these units to 50 separate test trials. Note increased jitter as latency increases.

contrast, the induction of synaptic-scaling is not dependent on NMDA receptors (Turrigiano et al. 1998) and is viewed as a different class of learning rules (Turrigiano and Nelson 2004). Thus mechanistically, the rule presented here is likely more similar to homeostatic plasticity.

The induction mechanisms of homeostatic plasticity are poorly understood (Turrigiano and Nelson 2004). However, it is clear that presynaptic-dependent scaling would require an additional mechanism to allow for the interaction between the long-term activity levels of both the pre- and postsynaptic terminals. Such a requirement could be implemented by retrograde or orthograde messengers implicated in other forms of plasticity (e.g., Chevaleyre and Castillo 2003; Sjöström et al. 2003). Alternatively, biochemical mechanisms could be present in the postsynaptic terminal that tracks the activity levels of the presynaptic terminal by integrating glutamate signals over the course of hours/days. Under this scenario, the implementation of the rule would still be local in relation to the postsynaptic neuron.

While presynaptic-dependent scaling has not been examined experimentally, it is consistent with experiments to date, because most studies that have characterized synaptic scaling were performed with global manipulations in which all presynaptic cells would be expected to have similar A_{pre} values. One clear prediction of the proposed learning rule is that, if only a subset of cells within a network are inactivated, postsynaptic cells would preferentially modify synapses from presynaptic cells that were not manipulated. In other words, neurons would preferentially strengthen synapses that would be more likely to increase the cell's activity level. While this type of experiment has not been explicitly examined, recent work has shown that local manipulation of postsynaptic activity does result in homeostatic synaptic plasticity (Burrone et al. 2002).

Synfire chains

The spatial-temporal pattern of activity produced by the network is similar to that of a synfire chain (Abeles 1991; Diesmann et al. 1999). Synfire chains have been proposed as a mode of activity propagation in neural networks; however, no learning rules have described how they may emerge in a recurrent network in a self-organizing manner. While presynaptic-dependent scaling does address this issue, it is important to note that the data presented here does not account for the hypothetical millisecond precision in synfire chains (Fig. 8). Indeed, consistent with recent experimental data (Buonomano 2003), the spike variability across trials increases as a function of time.

Spike-timing–dependent plasticity

Spike-timing–dependent plasticity (STDP) (Bi and Poo 1998; Feldman 2000; Markram et al. 1997) does not produce the type of dynamics observed here. Indeed, if STDP by itself is implemented in the preceding networks, the temporal pattern is abolished. For example, in Fig. 4*B*, STDP will produce a leftward shift of the spatial-temporal pattern as a result of long-term potentiation (LTP) of all the pre → post patterns. The leftward shift removes much of the temporal information and favors global synchronization, which leads to instability. However, it should be emphasized that STDP implemented together with presynaptic-dependent scaling can play a role in increasing the robustness and reliability of the responses in the

presence of noise (Supplemental Fig. 1). Specifically, by strengthening the synapses in the direction of propagation (pre → post), the trajectory can become more robust in response to perturbations. The problem is that STDP must be carefully balanced with presynaptic-dependent scaling. Thus future work will have to determine if this is a biologically plausible scenario.

Calculating the average activity level

All forms of homeostatic plasticity, including the one presented here, and some models of associative plasticity (Bienenstock et al. 1982) rely on the estimation of $A_i(t)$, which represents the average firing rate over a long time window. How neurons calculate $A_i(t)$ is not a trivial matter. For simplicity sake, let's assume a neuron has a target activity rate of 0.1 Hz: if a stimulus is presented every 10 s, a neuron must fire once per stimulus to achieve its setpoint (as in the preceding simulations). If a stimulus is presented every 5 s, it would have to fire 0.5 times per stimulus, whereas if the intertrial interval was 60 s, it would have to fire 6 times per stimulus (even if we allow for spontaneous activity similar problems arise). Clearly such a strong dependence on the ITI is not likely to be physiological.

While there is good evidence that cells do keep track of their average level of activity (Turrigiano and Nelson 2004; Turrigiano et al. 1998), the mechanisms remain unknown. Thus it is too early to rigorously address how the dependence on intertrial interval may be solved. However, the assumption that A_i would represent activity, not over absolute time, but over states in which the animal or network is attentive or behaving, as opposed to in rest or sleep states. Specifically, the presence of a salient stimulus would result in the beginning of an integration period that would last for some predetermined period of time.

How could such selective integration work? It has been proposed that cells may track their average activity through Ca^{2+} sensors with long integration times (Liu et al. 1998). In vivo selective integration over some states but not others could be achieved through the same ascending neuromodulatory mechanisms that gate plasticity (e.g., Bear and Singer 1982; Kilgard 2003; Kirkwood et al. 1999). Indeed, cortical levels of acetylcholine are different during sleep, quiet wakefulness, and active wakefulness (Marrosu et al. 1995). Thus a modulator such as acetylcholine could regulate either Ca^{2+} influx or the downstream integration of Ca^{2+} signals.

Weaknesses

One potential shortcoming of the model presented here is in relation to when more than one input is presented to the network during development. For example, what happens if on alternating trials two distinct groups of neurons are used as the input drive, and the target value of each cell remains at one? In contrast to the above simulations, there is not a unique solution to this scenario: a given neuron could fire twice on alternating trials or once on each trial. Simulations revealed that the network generally converges to one of these two solutions, and which solution was reached was dependent on initial parameters, including network connectivity and the magnitude of inhibition. However, I propose that cortical networks used for timing would indeed be triggered by a primary input and thus code time from the onset of the event whether it represented a single well-defined stimulus or a global multi-modal signal (Karmarkar and Buonomano 2003).

Timing

This model establishes how a recurrent network can implement a population clock in a self-organizing manner. While it is well established that temporal processing is a fundamental component of sensory and motor function, the neural mechanisms underlying temporal processing are not known (Ivry 1996; Mauk and Buonomano 2004). As mentioned above, it has recently been shown that subset of neurons in area HVc of the songbird implement a sparse code for time (Hahnloser et al. 2002). While the neural mechanisms by which this code is generated are not understood, this model provides a hypothesis as to how a sparse temporal code can emerge from a recurrent network.

ACKNOWLEDGMENTS

The author thanks T. Blair, V. Buonomano, J. Feldman, H. Johnson, U. Karmarkar, P. Latham, and C. Marder for helpful discussions and/or reading earlier versions of this manuscript.

GRANTS

This research was supported by National Institute of Mental Health Grant MH-60163.

REFERENCES

Abeles M. *Corticonics.* Cambridge, MA: Cambridge University Press, 1991.

Armstrong-James M, Diamond M, and Ebner FF. An innocuous bias in whisker use in adult rats modifies receptive fields of barrel cortex neurons. *J Neurosci* 14: 6978–6991, 1994.

Bear MF and Singer W. Modulation of visual cortical plasticity by acetylcholine and noradrenaline. *Nature* 320: 172–176, 1986.

Beaulieu Kisvarday Z, Somogyi P, Cynader M, and Cowey A. Quantitative distribution of GABA-immunopositive and -immunonegative neurons and synapses in the monkey striate cortex (area 17). *Cereb Cortex* 2: 295–309, 1992.

Beggs JM and Plenz D. Neuronal avalanches are diverse and precise activity patterns that are stable for many hours in cortical slice cultures. *J Neurosci* 24: 5216–5229, 2004.

Bi G-Q and Poo M-M. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J Neurosci* 18: 10464–10472, 1998.

Bienenstock EL, Cooper LN, and Munro PW. Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *J Neurosci* 2: 32–48, 1982.

Borst A and Theunissen FE. Information theory and neural coding. *Nat Neurosci* 2: 947–957, 1999.

Bottjer SW and Arnold AP. Developmental plasticity in neural circuits of a learned behavior. *Ann Rev Neurosci* 20: 459–481, 1997.

Brown TH, Kairiss EW, and Keenan CL. Hebbian synapses: biophysical mechanisms and algorithms. *Annu Rev Neurosci* 13: 475–511, 1990.

Buonomano DV. Decoding temporal information: a model based on short-term synaptic plasticity. *J Neurosci* 20: 1129–1141, 2000.

Buonomano DV. Timing of neural responses in cortical organotypic slices. *Proc Natl Acad Sci USA* 100: 4897–4902, 2003.

Burrone J, O'Byrne M, and Murthy VN. Multiple forms of synaptic plasticity triggered by selective suppression of activity in individual neurons. *Nature* 420: 414–418, 2002.

Chevaleyre V and Castillo PE. Heterosynaptic LTD of hippocampal GABAergic synapses: a novel role of endocannabinoids in regulating excitability. *Neuron* 38: 461–472, 2003.

Dayan P and Abbott LF. *Theoretical Neuroscience.* Cambridge, MA: MIT Press, 2001.

Destexhe A, Mainen ZF, and Sejnowski TJ. An efficient method for computing synaptic conductances based on a kinetic model of receptor binding. *Neural Comput* 6: 14–18, 1994.

Diesmann M, Gewaltig M-O, and Aertsen A. Stable propagation of synchronous spiking in cortical neural networks. *Nature* 402: 529–533, 1999.

Doupe AJ and Kuhl PK. Birdsong and human speech: common themes and mechanisms. *Ann Rev Neurosci* 22: 567–631, 1999.

Echevarría D and Albus K. Activity-dependent development of spontaneous bioelectric activity in organotypic cultures or rat occipital cortex. *Dev Br Res* 123: 151–164, 2000.

Fee MS, Kozhevnikov AA, and Hahnloser RHR. Neural mechanisms of vocal sequence generation in the songbird. *Ann NY Acad Sci* 1016: 153–170, 2004.

Feldman DE. Timing-based LTP and LTD at vertical inputs to Layer II/III pyramidal cells in rat barrel cortex. *Neuron* 27: 45–56, 2000.

Gawne TJ, Kjaer TW, Hertz JA, and Richmond BJ. Adjacent visual cortical complex cells share about 20% of their stimulus-related information. *Cereb Cortex* 6: 482–489, 1996.

Gupta A, Wang Y, and Markram H. Organizing principles for a diversity of GABAergic interneurons and synapses in the neocortex. *Science* 287: 273–278, 2000.

Hahnloser RHR, Kozhevnikov AA, and Fee MS. An ultra-sparse code underlies the generation of neural sequences in a songbird. *Nature* 419: 65–70, 2002.

He J, Hashikawa T, Ojima H, and Kinouchi Y. Temporal integration and duration tuning in the dorsal zone of the cat auditory cortex. *J Neurosci* 17: 2615–2625, 1997.

Hines ML and Carnevale NT. The NEURON simulation environment. *Neural Comput* 9: 1179–1209, 1997.

Ikegaya Y, Aaron G, Cossart R, Aronov D, Lampl I, Ferster D, and Yuste R. Synfire chains and cortical songs: temporal modules of cortical activity. *Science* 304: 559–564, 2004.

Ivry R. The representation of temporal information in perception and motor control. *Curr Opin Neurobiol* 6: 851–857, 1996.

Jimbo Y, Tateno T, and Robinson HPC. Simultaneous induction of pathway-specific potentiation and depression in networks of cortical neurons. *Biophys J* 76: 670–678, 1999.

Karmarkar UR and Buonomano DV. A model of spike-timing dependent plasticity: one or two coincidence detectors? *J Neurophysiol* 88: 507–513, 2002.

Karmarkar UR and Buonomano DV. Temporal specificity of perceptual learning in an auditory discrimination task. *Learn Mem* 10: 141–147, 2003.

Kilgard M. Cholinergic modulation of skill learning and plasticity. *Neuron* 38: 678–680, 2003.

Kilgard MP and Merzenich MM. Plasticity of temporal information processing in the primary auditory cortex. *Nat Neurosci* 1: 727–731, 1998.

Kirkwood A, Rozas C, Kirkwood J, Perez F, and Bear MF. Modulation of long-term synaptic depression in visual cortex by acetylcholine and norepinephrine. *J Neurosci* 19: 1599–1609, 1999.

Koch KW and Fuster JM. Unit activity in monkey parietal cortex related to haptic perception and temporary memory. *Exp Brain Res* 76: 292–306, 1989.

Laurent G, Wehr M, and Davidowitz H. Temporal representations of odors in an olfactory network. *J Neurosci* 16: 3837–3847, 1996.

Liu Z, Golowasch J, Marder E, and Abbott LF. A model neuron with activity-dependent conductances regulated by multiple calcium sensors. *J Neurosci* 18: 2309–2320, 1998.

Maass W and Sontag ED. Analog neural nets with gaussian or other common noise distributions cannot recognize arbitrary regular languages. *Neural Comput* 11: 771–782, 1999.

Marder CP and Buonomano DV. Timing and balance of inhibition enhance the effect of LTP on cell firing. *J Neurosci* 24: 8873–8884, 2004.

Markram H, Lubke J, Frotscher M, and Sakmann B. Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* 275: 213–215, 1997.

Markram H, Wang Y, and Tsodyks M. Differential signaling via the same axon of neocortical pyramidal neurons. *Proc Natl Acad Sci USA* 95: 5323–5328, 1998.

Marrosu F, Portas C, Mascia MS, Casu MA, Fa M, Giagheddu M, Imperato A, and Gessa GL. Microdialysis measurement of cortical and hippocampal acetylcholine release during sleep-wake cycle in freely moving cats. *Brain Res* 671: 329–332, 1995.

Mauk MD and Buonomano DV. The neural basis of temporal processing. *Annu Rev Neurosci* 27: 304–340, 2004.

Medina JF, Garcia KS, Nores WL, Taylor NM, and Mauk MD. Timing mechanisms in the cerebellum: testing predictions of a large-scale computer simulation. *J Neurosci* 20: 5516–5525, 2000.

Miller KD and MacKay DJC. The role of constraints in Hebbian learning. *Neural Comput* 6: 100–126, 1994.

Muller D, Buchs PA, and Stoppini L. Time course of synaptic development in hippocampal organotypic cultures. *Brain Res Dev Brain Res* 71: 93–100, 1993.

Pearlmutter BA. Gradient calculation for dynamic recurrent neural networks: a survey. *IEEE Trans Neural Netw* 6: 1212–1228, 1995.

Renart A, Song P, and Wang X-J. Robust spatial working memory through homeostatic synaptic scaling in heterogeneous cortical networks. *Neuron* 38: 473–485, 2003.

Ringach DL, Hawken MJ, and Shapley R. Dynamics of orientation tuning in macaque primary visual cortex. *Nature* 387: 281–284, 1997.

Rumelhart DE and McClelland JL. *Parallel Distributed Processing: Foundations.* Cambridge, MA: MIT Press, 1986.

Sjöström PJ, Turrigiano GG, and Nelson SB. Neocortical LTD via coincident activation of presynaptic NMDA and cannabinoid receptors. *Neuron* 39: 641–648, 2003.

Somers DC, Nelson SB, and Sur MJ. An emergent model of orientation selectivity in cat visual cortical simple cells. *J Neurosci* 269: 5448–5457, 1995.

Swadlow HA. Efferent neurons and suspected interneurons in S-1 vibrissa cortex of the awake rabbit: receptive fields and axonal properties. *J Neurophysiol* 62: 288–308, 1989.

Turrigiano GG, Leslie KR, Desai NS, Rutherford LC, and Nelson SB. Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature* 391: 892–896, 1998.

Turrigiano GG and Nelson SB. Homeostatic plasticity in the developing nervous system. *Nat Neurosci Rev* 5: 97–107, 2004.

van Rossum MCW, Bi GQ, and Turrigiano GG. Stable Hebbian learning from spike timing-dependent plasticity. *J Neurosci* 20: 8812–8821, 2000.

Weesberg J, Stambaugh CR, Kralik JD, Beck PD, Laubach M, Chapin JK, Kim J, Biggs SJ, Srinivasan MA, and Nicolelis MAL. Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature* 408: 361, 2000.