

Bistability and Tunable Gating in a Recurrent Neuronal Network with Synaptic Failure and Hidden Neurons

Louis Tao, David Cai, and Michael J. Shelley

Courant Institute of Mathematical Sciences, New York University, 251 Mercer Street, New York, NY 10027

(Dated: September 19, 2002)

We study networks of all-to-all coupled, integrate-and-fire, excitatory neurons, with a portion of the population receiving feedforward drive and the remaining, “hidden”, portion receiving no feedforward input. We demonstrate that this system undergoes a subcritical bifurcation as either the input or recurrent coupling is varied. Because of this transition, the network gates feedforward inputs and exhibits bistability and hysteresis. The hidden population allows this network response to be continuously tunable in the forcing and coupling strengths. Furthermore, these features persist in the presence of synaptic failure and for small network sizes, suggesting that the computations discussed here could be implemented in biological networks. Finally, we demonstrate that a long network correlation, orders of magnitude longer than the synaptic time, emerges from the recurrent coupling. This correlation time scales with network size and with synaptic connection probability.

While the disordered topology of networks in social, ecological and biological systems is much studied [1, 2], the understanding of dynamical processes on complex networks is just beginning. Issues such as network tolerance, robustness and structural stability [3, 4] have been considered in recent studies of the modeling of dynamics of networks with fixed, though random, topology [5, 6]; examples include synchronization on small-world [7] and scale-free networks [8]. However, many biological systems are characterized by topologies that change on time-scales comparable to or even faster than the dynamical time-scales of their constitutive elements. We note that of the many processes that regulate synapses on a variety of different time-scales, e.g., involved in short-term synaptic depression, long-term potentiation, and plasticity [9, 10], the connectivity changes induced by synaptic failure can be viewed as occurring on time-scales much faster than the synaptic time-scales. Whether these topological changes have biological functions and what kind of computations these networks are capable of are important and unresolved issues.

In the neuroscience literature, much theoretical attention has focused on the role of excitation of single population, strongly coupled networks [11–14], or on feedforward, hierarchical models [15, 16]. Although it has been shown that recurrent coupling can have a significant role in network oscillation and synchrony [12, 14], the possible cortical functions of this coupling has yet to be fully clarified.

Here we show that a simple modification of an all-to-all

coupled network of excitatory, integrate-and-fire (I&F) neurons leads to a number of unexpected results. Our network differs from other studies in a single respect: only a portion of the neurons in the network receives feedforward excitation. The “hidden” remainder receives no feedforward input. A prime motivation behind our model comes from the observation that there are many neurons without direct, lower-level, input in even the earliest cortical areas [17]. In the mammalian primary visual cortex, for instance, neurons outside of the so-called input layers receive no direct drive from the thalamus. It has been long believed that such neurons are part of a hierarchy of early visual processing, where “Simple” cells receive visual drive and “Complex” cells are driven by the summed outputs of the Simple cells [18]. In recent revisions of this picture, Simple and Complex cells arise in a single recurrent network, where those neurons without feedforward input are driven entirely by forcings pooling Simple and Complex cells [19–21]. Here we use an idealization of such network models to explore the possible biological functions of a subpopulation of neurons, i.e., hidden (Complex), driven entirely and recurrently by network inputs.

The dynamical properties of our model illustrate the importance of network-based mechanisms. Other recent work has utilized recurrent networks to model persistent activity and stable integration in various models of working memory, oculomotor control and decision making [22–24]. However, these models require individual neurons with long synaptic time-scales [35] or other biophysically complex properties [24, 25]. In contrast, similar computations achieved by our network do not rely on such special properties, but utilizes instead the structured recurrent coupling between neurons within the network.

We demonstrate that this network can act as a *gate* for the feedforward drive: The hidden cells fire only when the input drive is above a certain threshold, and when they fire, they fire at a high rate. This particular phenomenon is continuously tunable by synaptic strength, synaptic failure probability and input drive and is robust to the variations in these quantities. This “phase” transition is subcritical and the network becomes bistable and hysteretic in either one of two parameters, the strength of the input and the strength of the network coupling. Furthermore, in the presence of synaptic failure, the network state is asynchronous (in contrast to the case of no failure, where the hidden subpopulation is synchronous).

The observation that the subcritical transition is robust in the presence of input noise, inhibitory coupling, synaptic failure and persist in small networks suggests that this gating mechanism may be implemented biologically. It is conceivable that graded global computations of high complexity can be achieved by grouping together a hierarchy of such networks spatially distributed with different synaptic strengths or failure rates.

The use of synaptic failure to modify synaptic strength and network dynamics has received little attention. Based on our results below, we speculate that synaptic failure of biological networks could be usefully regulated by higher order processes so as to control the dynamics and the transfer of information [26] on certain time scales. These time-scales may be different from those involved in the strengthening of individual synapses, e.g., processes thought to be underlying development or learning [27].

The Model. Our model consists of an all-to-all coupled network of N excitatory I&F, conductance-based point neurons. Individual membrane potentials, $V_i(t)$, follow

$$\frac{dV_i}{dt} = -g_R(V_i - V_R) - g_i^E(V_i - V_E), \quad i = 1, \dots, N \quad (1)$$

where g_R is the leak conductance; g_i^E is the excitatory conductance of the i^{th} neuron; V_R and V_E are the reset and excitatory reversal potentials, respectively. In an I&F neuron, when the membrane potential V_i hits a threshold, V_T , that time is recorded as a spike time, and V_i is instantaneously reset to the reset potential. Action potentials are then delivered across the network, inducing conductance changes in other neurons. After each spike, V_i is held at rest for an absolute refractory period. (In our results reported below, we typically use a refractory period $\tau_{\text{ref}} = 3$ ms, consistent with physiological measurements [9].) Eq. (1) is in non-dimensional form with $V_T = 1$, $V_R = 0$, and, by using the commonly accepted values [9] for the various biophysical parameters, $V_E = 4.67$ and $g_R = 50 \text{ sec}^{-1}$. Only time retains dimension. To simulate Eq. (1) numerically, we modified a second-order Runge-Kutta time-stepping scheme to treat accurately the discontinuities in the membrane potential due to spike resets [28, 29].

The network model is closed by the specification of the individual g_i^E 's:

$$g_i^E(t) = \frac{S_E}{N} \sum_j \sum_k P_i^{jk} G(t - t_j^k) + f_i(t), \quad i = 1, \dots, N \quad (2)$$

where t_j^k is the k^{th} spike-time of the j^{th} neuron. $G(t)$ is the postsynaptic conductance (PSC), induced by an action potential. Typically, only half of the neurons are driven by feedforward excitation in the form of action potentials: $f_i(t) = c_i \sum_k G(t - s_i^k)$, for $i = 1, \dots, N/2$ [$f_i(t) = 0$, otherwise]. The k^{th} "spike-time" for the i^{th} neuron, s_i^k , is given by a Poisson process




FIG. 1: Network Gating. Firing rates as a function of g_{input} and pS_E (fixing $S_E = 1$). Solid: hidden population firing rate; dotted: driven population firing rate. The dark dashed line indicates the gating threshold as a function of pS_E and g_{input} . For each g_{input} , the firing rates are obtained by increasing p . See text for hysteretic effects.

with rate ν_e . Here we take $c_i = 0.05$. We model each PSC by a difference of exponentials: $G(\tau) = \frac{1}{\tau_1 - \tau_2} (\exp(-\tau/\tau_1) - \exp(-\tau/\tau_2)) \theta(\tau)$, where θ is the Heaviside function and τ is the time after a spike. We typically set $\tau_1 = 5$ ms and $\tau_2 = 1$ ms, which are values used to model fast excitatory synaptic couplings [36]. The strength of the input drive can be adjusted by changing ν_e . The average conductance induced by the feedforward forcing is then $g_{\text{input}} = c_i \nu_e$, since $G(t)$ has integral one.

Another distinguishing ingredient of our model is synaptic failure. Although it is observed that synaptic transmission becomes more reliable as the temperature is increased, a recent study using *in vitro* preparations of rat visual cortex found that, even at a physiological temperature of 36°C , the mean failure probability can be more than 10% [30].

We model synaptic failure by the random variable P_i^{jk} . In traditional network modeling P_i^{jk} is implicitly fixed to unity, while here P_i^{jk} is chosen to be 1 with probability p . At each spike-time t_j^k , each postsynaptic neuron i has a probability p of responding to an action potential from presynaptic neuron j . Thus $p = 1$ is no synaptic failure, and, in general, $1 - p$ is the probability of synaptic failure. We assume that p is independent of time. In Eq. (2), since $G(t)$ has integral one and N normalizes the double sum to order 1, the product pS_E represents an effective synaptic strength.

Gating and Bistability. Figure 1 displays the population firing rate as a function of g_{input} and pS_E in a $N = 1024$ network. The network exhibits a gating effect in pS_E , namely, for each level of feedforward input, there is a critical coupling strength below which the hidden cells are silent and above which the hidden cells are firing




FIG. 2: Finite-size effects on gating. At $g_{\text{input}} = 10 \text{ sec}^{-1}$ and $S_E = 1$, the firing rate of the hidden population in differently sized networks.




FIG. 3: Effect of refractory period. The firing rate of the hidden population in a $N = 1024$ network, $g_{\text{input}} = 30 \text{ sec}^{-1}$, as a function of pS_E , for different τ_{ref} .

vigorously. This gating response is also seen when pS_E is fixed and g_{input} is varied. Hence, the critical threshold can be tuned continuously (see thick, dashed line in Fig. 1). For instance, by choosing (and fixing) pS_E appropriately, we can set the threshold for the input drive.

We find that gating also occurs in much smaller networks. Figure 2 shows the firing rate of hidden cells, as a function of increasing p , in networks of different sizes. The two largest networks are indistinguishable: At $N = 256$, the firing rate curve has basically converged to the large N limit. However, even in a network of 16 neurons, the jump in the firing rate remains sharp. This persistence to small N ($\sim 10-20$) networks is noteworthy in that a small group of neurons can be quickly recruited to perform necessary computations.

One effect of the absolute refractory period is to pre-




FIG. 4: Bifurcation and hysteresis in a $N = 1024$ network. The parameters are $g_{\text{input}} = 15 \text{ sec}^{-1}$ and $S_E = 1$. The population firing rate in thick lines. The jump in the population firing rate occurs at different pS_E depending on whether p is increased from $p = 0.25$ or decreased from $p = 0.45$. Arrows indicate the direction that p is changed. Solutions from our mean-field model in thin lines (solid and dashed indicate stable and unstable solutions respectively).

vent runaway network amplification. Figure 3 shows the hidden population firing rate for three choices of the refractory period, τ_{ref} , as a function of pS_E . A sharp transition occurs in all cases. However, the choice of $\tau_{\text{ref}} = 3 \text{ ms}$ permits the firing rates to saturate at a rate consistent with physiological measurements.

The firing rates in Figs. 1-3 were obtained by increasing p (at fixed $S_E = 1$) for each g_{input} . The sharp jumps in the population firing rate as pS_E is increased indicate a subcritical bifurcation, where bistability and hysteresis are usually observed. This is indeed the case for our network. Figure 4 shows the two different firing rate curves obtained by increasing or decreasing p . The network gates when we increase p above a critical p and the hidden cells become active. If we then decrease p , the hidden cells remain active below this first critical p for an extended range in p .

The observed network hysteresis is well captured by the solution of a mean-field analysis (see forthcoming note), which is also displayed in Fig. 4. Our analysis suggests that this transition is different than the typical subcritical case, where the lower solution loses stability at the bifurcation point. Here the lower solution — driven cells firing and hidden cells silent — disappears at the critical pS_E where the hidden neurons start to fire.

Figure 5a shows the firing rate curves for networks of different sizes. Bistability and hysteresis in pS_E persist for small networks (e.g., $N = 20$), though the range of bistability is reduced relative to large networks. Similarly, the network bistability in g_{input} persists in small networks. This is summarized in Fig. 5b, which displays




FIG. 5: Hysteresis in small networks. In both figures, the solid line is the mean firing rate of the hidden population from a $N = 1024$ network and the dashed line is from a $N = 20$ network. (a) Hysteresis as p is varied ($g_{\text{input}} = 10 \text{ sec}^{-1}$ and $S_E = 1$). (b) Hysteresis as g_{input} is varied ($p = 0.45$, $S_E = 1$).

the firing rate curves as g_{input} is varied. In this particular case, it is possible to keep the hidden population active even as g_{input} is decreased to zero. This is an instance of persistent activity [22]. In this case, once the hidden neurons are turned on, it will take an inhibitory current to turn them off. This persistent activity and hysteretic behavior can be used to construct integrator circuits [24].

While finite-size effects in small networks are not easily accounted for theoretically, our simulations show that if a synaptic time, τ_1 , is increased from 5 ms to 80 ms [37], then even a network of 4 neurons (2 hidden and 2 driven) can remain firing as $g_{\text{input}} \rightarrow 0^+$. Without a long synaptic time constant, however, the hidden cells in this 4-neuron network become silent before $g_{\text{input}} = 0$. For this case, since our mean-field solutions (which assume $N \gg 1$) have a finite firing rate as $g_{\text{input}} \rightarrow 0^+$, finite-size effects may be less significant in systems with longer excitatory time-courses.

Subcritical transitions in all-to-all networks have been uncovered previously [22]; however, there are some important features that arise only in the presence of the hidden population. In conventional all-to-all networks, where all cells are driven, the network threshold at which a subcritical bifurcation occurs as the input drive is increased is the same as the single neuron spike threshold. This threshold is independent of the coupling strength (although it can be changed through current injection or by adjusting the leak conductance [24]). With the incorporation of hidden cells, a second subcritical bifurcation arises, leading to a threshold different from that of the directly driven neurons. Consequently, this new threshold is continuously tunable in both g_{input} and $p S_E$: The threshold in g_{input} can be varied by adjusting the network $p S_E$ and vice versa. The fact that the threshold depends




FIG. 6: Dynamics of Single Neurons and Collective Network Dynamics. (a) The membrane potential $V_i(t)$ and (b) excitatory conductance $g_i^E(t)$ of a hidden neuron in a network with synaptic failure, $p = 0.99$, $S_E = 0.5051$, and $g_{\text{input}} = 10 \text{ sec}^{-1}$; (d) $V_i(t)$ and (e) $g_i^E(t)$ of a hidden neuron in a network without synaptic failure, $p = 1.00$, $S_E = 0.5$, and $g_{\text{input}} = 10$. (c) Spike rasters for the $p = 0.99, S_E = 0.5051$ network; (f) Spike rasters for the $p = 1.00, S_E = 0.5$ network. All simulations were performed with $N = 1024$. For clarity, in (c) and (f), the rasters are only shown for the 64 neurons indexed between 480 and 543 (index not labeled). Here neurons with $i > 512$ (upper half of the rasters) are of the hidden type.

continuously of both g_{input} and $p S_E$ permits the gating response of this network to be robust to noise in the feed-forward drive and to fluctuations in network parameters. That is, the network need not be precisely tuned to function as a gate. Furthermore, the network bifurcation operates in a regime where each individual driven neuron is sufficiently forced to fire regardless of the history of the network state. Thus the network behavior is dynamically reproducible and easily controlled.

Effects of Synaptic Failure. While Fig. 1 was produced by fixing $S_E = 1$ and increasing p , the diagram remains qualitatively the same for fixed $p = 1$ and increasing S_E : Synaptic failure does not alter the gating feature of this network. Even though the firing rate diagrams of the fixed $S = 1$ and the fixed $p = 1$ networks are quantitatively similar, the collective *spike dynamics* of the two networks are very different. Figures 6a,d show $V(t)$ of two typical hidden neurons, in networks with and without failing synapses, respectively. It is hard to distinguish between the two networks on the basis of individual membrane potentials alone. However, examination of the dynamics of the excitatory conductances of the same two neurons, displayed in Figs. 6b,e, reveals that the network without synaptic failure is oscillating. This oscillation is a consequence of the synchrony in the hidden subpopulation, as seen in the rasters of spike-times shown in Fig. 6f

TABLE I: Role of Synaptic Failure.

$p =$	0.0 ^a	0.1	0.5	0.99	1.0	1.0 ^b
$CV_{output,1}$ ^c	0.036	0.046	0.036	0.035	0.056	0.034
$CV_{output,2}$ ^d	0.015	0.037	0.015	0.009	0.014	0.021
CV_{input}	0.999	0.998	0.997	1.003	1.886	0.995
$\sigma(V_1)$ ^e	0.794	0.792	0.793	0.793	0.774	0.795
$\sigma(V_2)$ ^f	0.779	0.770	0.778	0.778	0.0	0.784

^aThe $p = 0$ column is obtained in a uncoupled network, with the feedforward inputs into both populations matched to the total inputs into the two populations of the $p = 0.99$ simulation.

^bIn this set of $p = 1$ simulations, the hidden cells receive *weak* feedforward input $g_{input} = 0.5 \text{ sec}^{-1}$ (compared to $g_{input} = 10 \text{ sec}^{-1}$ for the driven population).

^cCV of the output spikes of the driven neurons

^dCV of the output spikes of the hidden neurons

^eNormalized standard deviation $\sigma(V)$ of the driven neurons (see text)

^f $\sigma(V)$ of the hidden neurons

for the $p = 1$ simulation. In this case, since there is no synaptic failure, individual hidden neurons have identical $g_i^E(t)$ and the hidden population becomes synchronized after a short transient (not shown). However, a little synaptic unreliability will desynchronize this population. Figure 6c shows spike rasters for a $p = 0.99$ simulation. Already at a synaptic failure probability of 1%, the network is asynchronous.

The differences between these two networks are also apparent when we examine the coefficient of variation (CV) of the spike times. CV is a measure of the variability of neuronal spike trains [9, 10, 31]; $CV = 1$ for Poisson spike trains and $CV = 0$ for periodic spike trains. Table I shows the CV for simulations of networks with the same synaptic coupling strength pS_E , but with different synaptic connection probability p . We present the CV of both the *output* and the *input* spike trains of individual neurons. The output CVs are computed by using the interspike interval (ISI) of individual neuronal spike trains and averaging over the driven and the hidden populations separately. The input CV describes the spike train that an individual neuron receives. Here we take the ISI of the entire set of network spike-times.

Synaptic failure has a significant effect on network synchrony. In the $p = 1.0$ case, the hidden neurons become synchronized and this leads to input spike-times that are not Poisson ($CV = 1.886 \neq 1$, in Table I). This synchronization can be removed by introducing noise either by injecting a weak feedforward input into the hidden neurons (see the second $p = 1.0$ column) or by introducing unreliable synapses. In either case, the input spike-times become Poisson ($CV \approx 1$). In all cases, the output CVs are small. This is due to the almost regular firing of individual neurons. The fact that the driven cells have an output CV larger than the hidden cells reflects the feedforward input received by the driven cells.

Table I also tabulates $\sigma(V) = (\langle V^2 \rangle - \langle V \rangle^2) / \langle V^2 \rangle$, the

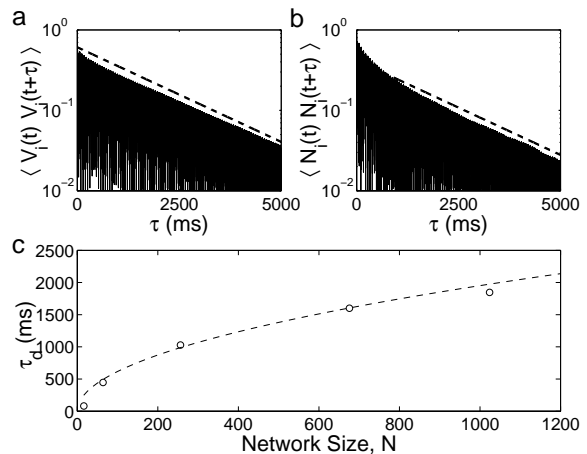


FIG. 7: Correlation time-scales and finite-size scaling. (a) $\langle V_i(t)V_i(t + \tau) \rangle$ of the hidden population for a $N = 1024$ network. $g_{input} = 10 \text{ sec}^{-1}$, $p = 0.9$, $S_E = 0.5556$. The dashes are an exponential fit. (b) $\langle N_i(t)N_i(t + \tau) \rangle$ of the hidden population in this network. The dashed line has the same slope as that in (a). The fine oscillations in the correlations are associated with the periodic firing of individual neurons. (c) The decay time of $\langle V_i(t)V_i(t + \tau) \rangle$ as a function of network size. The dashed line is proportional to the square root of the network size.

normalized standard deviation of the membrane potential for each of the two populations separately. ($\langle \dots \rangle$ denotes average over time and over subpopulations.) $\sigma(V)$ attains the value 0 when the population is totally synchronized. Indeed, the synchronized hidden cells in the noise-less $p = 1.0$ network possess $\sigma(V) = 0$. Furthermore, $\sigma(V) (= 0.770 - 0.795)$ of all other networks are virtually the same, demonstrating that synaptic failure not only desynchronizes the hidden population, but also allows the dynamics of these different asynchronous networks to be statistically similar.

Emergent time-scales. A long time-scale emerges from our recurrent network dynamics. Figure 7a shows the autocorrelation of the membrane potentials of the hidden population, $\langle V_i(t)V_i(t + \tau) \rangle$, as a function of time τ . The autocorrelation shows an exponential decay with a time constant, τ_d , of approximately 2 seconds. This time-scale is orders of magnitude longer than the synaptic time-scale of 5 ms or the leakage time-scale of 20 ms.

The same decay time also manifests itself in the correlation of the number of “connections”: $\langle N_i(t)N_i(t + \tau) \rangle$, where N_i denotes the number of postsynaptic targets after spikes of the i^{th} neuron. We show a coarse-grained version of this autocorrelation in Fig. 7b. We bin the number of postsynaptic connections to 1 ms (coarse-graining) windows and use this averaged $N_i(t)$ to estimate the correlation. The autocorrelations of V_i and N_i are related since a larger V_i implies that the neuron is closer to threshold, and is more likely to fire and thus secure more postsynaptic targets.

Figure 7c shows the decay time of these autocorrelations as a function of network size. This time-scale is longer in larger networks and in networks with a higher connection probability p (for fixed pS_E ; not shown) and scales roughly as the square root of the network size, reminiscent of the critical slow-down near certain phase transitions [32]. This scaling suggests that this network can establish its own time-scales, independent of the synaptic ones, by recruiting an appropriate number of active neurons. Thus synaptic failure may have an important role in associate memory and on-line learning [10].

In the results above, we have restricted discussion to the case of excitatory coupling and where the hidden cells do not receive feedforward drive. The gating and hysteretic features remain even when the hidden cells receive some feedforward input and when inhibitory neurons are included. Our mean-field analysis of the network including these additional effects suggests that the bifurcation is subcritical as long as the driven and the hidden populations receive different levels of feedforward inputs—and the larger this difference, the larger the region of bistability. The main effect of the input drive or of inhibition is to change some of the quantitative aspects, e.g., population firing rates after crossing the gating threshold, and size of the hysteretic region.

Furthermore, a long time-scale naturally arises from the recurrent coupling within our network. While network time-scales cannot be easily observed in experiments, there is some indirect evidence that Complex (hidden) cells may be responsible for long time-scales. We point out that in reverse correlation experiments of the macaque primary visual cortex, longer dynamical time-scales are typically found in neurons within layers 5 and 4B [33], which are also layers with large percentages of Complex cells [34].

While many of our model results can be formalized in a mean-field setting and made precise there, here we have given a phenomenological account, and relegate theoretical analyses to a forthcoming note. Our main purpose here is to illustrate that in a simple, finite-size setting, tunable and robust dynamics such as gating and bistability can be achieved by incorporating hidden cells and synaptic failure.

We thank Robert Shapley for useful conversations. D. Cai is supported by NSF DMS0211655 and a Sloan Fellowship. L. Tao is supported by NEI grant EY07158 and the Sloan-Swartz Theoretical Neurobiology Program.

-
- [1] Erdős, P & Rényi, A. (1959) *Publ. Math. (Debrecen)* **6**, 290.
 [2] Bollobas, B. (2001) *Random Graphs*. (Cambridge University Press, New York).
 [3] Strogatz, S. H. (2001) *Nature* **410**, 268–276.

- [4] Albert, R & Barabasi, A.-L. (2002) *Rev. Mod. Phys.* **74**, 47–97.
 [5] Golomb, D & Hansel, D. (2000) *Neural Comp.* **12**, 1095–1139.
 [6] Guardiola, X, Diaz-Guilera, A, Llas, M, & Perez, C. J. (2000) *Phys. Rev. E* **62**, 65565–5570.
 [7] Lago-Fernandez, L. F, Huerta, R, Corbacho, F, & Siguenza, J. A. (2000) *Phys. Rev. Lett.* **84**, 2758–2761.
 [8] Wang, X. F & Chen, G. (2002) *IEEE Trans. Circ. Syst. I* **42**, 430–447.
 [9] Koch, C. (1999) *Biophysics of Computation*. (Oxford University Press, Oxford).
 [10] Dayan, P & Abbott, L. (2001) *Theoretical Neuroscience*. (MIT Press, Cambridge).
 [11] Gerstner, W. (1995) *Phys. Rev. E* **51**, 738–758.
 [12] Brunel, N & Hakim, V. (1999) *Neural Comp.* **11**, 1621–1671.
 [13] Bressloff, P & Coombes, S. (2000) *Neural Comp.* **12**, 91–129.
 [14] Neltner, L, Hansel, D, Mato, G, & Meunier, C. (2000) *Neural Comp.* **10**, 467–483.
 [15] Diesmann, M, Gewaltig, M.-O, & Aertsen, A. (1999) *Nature* **402**, 529–533.
 [16] van Rossum, M. C. W, Turrigiano, G. G, & Nelson, S. B. (2002) *J. Neurosci.* **22**, 1956–1966.
 [17] Tanaka, K. (1985) *Vis. Res.* pp. 357–364.
 [18] Hubel, D & Wiesel, T. (1962) *J. Physiol (Lond)* **160**, 106–154.
 [19] Chance, F, Nelson, S, & Abbott, L. (1999) *Nature Neurosci.* **2**, 277–282.
 [20] Wielaard, J, Shelley, M, Shapley, R, & McLaughlin, D. (2001) *J. Neuroscience* **21**.
 [21] Tao, L, Shelley, M, McLaughlin, D, & Shapley, R. (2002) *submitted to Nature Neurosci.*
 [22] Wang, X.-J. (1999) *J. Neurosci.* **19**, 9587–9603.
 [23] Laing, C & Chow, C. (2001) *Neural Comp.* **13**, 1473–1494.
 [24] Koulakov, A, Raghavachari, S, Kepecs, A, & Lisman, J. (2002) *Nature Neurosci.* **5**, 775–782.
 [25] Lisman, J, Fellous, J.-M, & Wang, X.-J. (1998) *Nature Neurosci.* **1**, 273–275.
 [26] Faure, P, Kaplan, D, & Korn, H. (2000) *J. Neurophysiol.* **84**, 3010–2025.
 [27] Bi, G.-Q & Poo, M.-M. (2001) *Ann. Rev. Neurosci.* **24**, 139–166.
 [28] Hansel, D, Mato, G, Meunier, C, & Neltner, L. (2000) *Neural Comp.* **12**, 1607–1641.
 [29] Shelley, M & Tao, L. (2001) *J. Comput. Neurosci.* **11**, 111–119.
 [30] Hardingham, N. R & Larkman, A. U. (1998) *J. Physiol.* **507**, 249–256.
 [31] Koch, C & Segev, I. (1998) *Methods in Neural Modeling from Ions to Networks, 2nd edition*. (MIT Press, Boston).
 [32] Ma, S. K. (1976) *Modern Theory of Critical Phenomena*. (W. A. Benjamin, Reading, Massachusetts).
 [33] Pugh, M, Ringach, D, Shapley, R, & Shelley, M. (2000) *J. Comput. Neurosci.* **8**, 143–159.
 [34] Ringach, D, Shapley, R, & Hawken, M. (2002) *J. Neuroscience* **22**, 5639–5651.
 [35] mediated by, e.g., NMDA (N-methyl-D-aspartate) receptors; see, e.g., [9]
 [36] These time-scales are typical of AMPA; cf. [9].
 [37] NMDA is typically modeled by a synaptic time constant greater than 80 ms; cf. [9]