

Content-addressable memory with spiking neurons

R. Mueller and A. V. M. Herz

Innovationskolleg Theoretische Biologie, Humboldt-Universität zu Berlin, Invalidenstrasse 43, 10115 Berlin, Germany

(Received 25 June 1998)

Time evolution and cooperative phenomena of networks with spiking model neurons are analyzed with emphasis on systems that could be used as content-addressable memories. Stored memories are represented by distributed patterns of neural activity where single cells either fire periodically or remain quiescent. Two distinct mechanisms to generate relaxation behavior toward such periodic solutions are investigated: delayed feedback and subthreshold oscillations. Using theoretical analysis and numerical simulations it is shown that in both cases model networks with integrate-and-fire neurons possess storage capabilities similar to those of conventional associative neural networks. [S1063-651X(99)09202-8]

PACS number(s): 87.10.+e, 05.20.-y, 07.05.Mh, 64.60.Cn

I. INTRODUCTION

Recurrent neural networks may be programmed to function as content-addressable memories that recover stored patterns from incomplete or noisy inputs. To do so, correlations within patterns to be memorized are encoded in the synaptic weights. By this procedure, multiple patterns can be implemented as fixed-point attractors of the network dynamics. Starting from an initial state close to one of the stored attractors, the system dynamics relaxes to this attractor and thus retrieves the stored pattern [1–3].

Various brain regions have been hypothesized to operate as content-addressable memories, for example, the CA3 region in the hippocampus and locally connected systems of pyramidal cells in neocortical association areas [4–6]. Traditionally, these systems have been modeled using coarse-grained dynamical descriptions based on short-time averaged firing rates. This approach leads either to models with continuous time and real-valued state variables (“graded-response neurons”) [2,3] or to systems with discrete time and binary [1,7] or real-valued state variables [8,9]. Extensive theoretical results regarding the convergence properties and storage capacity of such “attractor neural networks” have been derived [10–12].

Most cortical neurons communicate using discrete pulses of electrical activity, called “action potentials” or “spikes.” Firing-rate models neglect this characteristic feature of neural signals. It is, therefore, important to compare the collective properties of biologically more realistic approaches with those of the traditional network models. Previous studies have demonstrated that systems with spiking model neurons offer computational possibilities not shared by models based on firing rates such as computations based on relative spike times [13–18]. On the other hand, it has also been shown that for stationary solutions, firing-rate descriptions cover a large class of networks with spiking model neurons [19].

The apparent discrepancy between these two lines of research is resolved if one realizes that results based on the dynamics of specific models can only demonstrate the richness of phenomena generated by that very class of model. Other models may or may not exhibit the same phenomena. Results obtained with specific models are nevertheless extremely useful to prove the feasibility of computational

mechanisms. It is in this spirit that we investigate associative capabilities of model networks with integrate-and-fire neurons.

This study focuses on systems in which stored binary memories are represented by distributed patterns of neural activity where single cells either fire periodically (the “on”-state) or remain completely quiescent (the “off”-state). Two different mechanisms to achieve such behavior are investigated in detail: delayed feedback within the model network and externally generated subthreshold oscillations. We demonstrate that both mechanisms give rise to collective properties that are almost identical with those of the Little model [7], one of the classical attractor neural networks. Since periodic membrane oscillations due to rhythmic background activity are typical for various brain regions, our results indicate that such oscillations may play a beneficial role in content-addressable memory processes [20]. The results also show that although the microscopic time evolution of integrate-and-fire neurons strongly differs from that of graded-response neurons or binary neurons, the collective network dynamics may nevertheless exhibit rather similar phenomena.

II. MODEL SYSTEMS

Biological neurons generate an action potential when their cell body (soma) is sufficiently depolarized. The action potential then propagates along the axon to synapses on the dendritic trees of downstream (postsynaptic) neurons. When the action potential arrives at a synapse it initiates the release of neurotransmitter which leads to a flow of ionic currents that depolarize or hyperpolarize the postsynaptic cell. Depending on the integrated inputs from the many thousand cells a cortical neuron is typically connected with, the postsynaptic neuron will in turn fire an action potential at some later time and thus influence further neurons. Feedback through recurrent connections results in complicated spatiotemporal dynamics.

Model neural networks are constructed to capture important features of these intricate dynamical processes. In order to allow for an analysis of the collective behavior of large feedback networks, the microscopic dynamics has to be simplified as much as possible. Individual neurons are therefore

often considered to be electrotonically compact. Within this class of model, “integrate-and-fire” neurons [16,21–29] are characterized by a particularly simple description in that the state of each neuron i ($1 \leq i \leq N$) is modeled by a single dynamical variable u_i , the membrane potential at the soma.

If the membrane potential u_i is below the firing threshold u_{thresh} , integrate-and-fire neurons operate as leaky integrators,

$$C \frac{du_i(t)}{dt} = -\frac{1}{R} [u_i(t) - u_0] + I_i(t). \quad (1)$$

The capacitance C and the resistance R of the cell membrane determine the membrane time constant $\tau_{RC} = RC$. In the absence of an input current $I_i(t)$, the membrane potential $u_i(t)$ approaches its rest value u_0 .

Within the mathematical formulation, the term u_0/R can be absorbed in $I_i(t)$ and we will therefore focus on the case $u_0 = 0$ without loss of generality. By rescaling time t and input current I_i , the capacitance C and resistance R can be taken as one and we arrive at

$$\frac{du_i(t)}{dt} = -u_i(t) + I_i(t). \quad (2)$$

When the membrane potential of an integrate-and-fire model neuron reaches the firing threshold u_{thresh} , the cell generates a uniform action potential modeled as a δ pulse, and the membrane potential is reset to $u_{\text{reset}} < u_{\text{thresh}}$. The output of an integrate-and-fire neuron is thus just a sequence of δ pulses.

For convenience, units can be chosen such that $u_{\text{thresh}} = 1$ and $u_{\text{reset}} = 0$. Since the reset is assumed to be instantaneous, the membrane potential u_i takes two values when cell i generates an action potential at time t . Where necessary, these two values will be denoted by $u_i(t^-)$ and $u_i(t^+)$. To simplify the notation, temporal arguments otherwise always refer to the time immediately prior to firing. Special care has to be taken in systems with δ -shaped input currents that may cause a cell to generate an action potential at time t even if immediately prior to that instant, its membrane potential is strictly below the firing threshold. In particular, a positive input of the form $I \delta(t - \tilde{t})$ will cause an action potential in a subthreshold cell if $I \geq 1 - u_i(\tilde{t}^-)$. This example shows that in models with integrate-and-fire neurons one has to take the dynamical consequences of δ -shaped postsynaptic currents into careful consideration.

The reset process can be formally included in the time evolution (2) by an additional effective current I_i^{reset} that guarantees that the membrane potential u_i is reset to zero immediately after neuron i generates an action potential. Denoting input currents to cell i from other neurons in the modeled network by I_i^{network} and input currents due to external stimuli and background activity by I_i^{external} , we get

$$I_i(t) = I_i^{\text{network}}(t) + I_i^{\text{external}}(t) + I_i^{\text{reset}}(t), \quad (3)$$

where

$$I_i^{\text{reset}}(t) = -\sum_f u_i(t_{i,f}^-) \delta(t - t_{i,f}). \quad (4)$$

Here, $t_{i,f}$ with $f \in \mathbb{N}$ are the times when neuron i generates an action potential. The factor $u_i(t_{i,f}^-)$ takes the δ contributions mentioned above into account and guarantees that the membrane potential of cell i is reset to zero.

From a biological point of view, Eq. (4) describes a situation where ionic currents due to spike generation are so strong that they override any simultaneous input currents. Consistent with this picture, currents caused by previous synaptic inputs will be set to zero when a cell emits an action potential. Backpropagating action potentials [30] could provide a biophysical mechanism for this phenomenon in biological neurons. In order to properly formulate this characteristic of the dynamics, let us now focus on the postsynaptic effects of spike activity.

If, say, neuron j spikes, an action potential travels along its axon to other neurons. When it arrives at a synapse with neuron i , neurotransmitter is released and triggers a postsynaptic current in cell i . The shape of the current resulting at the soma is denoted by $\alpha(\tau)$ where τ measures the time since the action potential arrived at the synapse. The functional form of $\alpha(\tau)$ may thus be used to describe the effects of synaptic transmission and/or passive dendritic conduction, with $\alpha(\tau) = 0$ for $\tau < 0$. The input to cell i from other neurons is then

$$I_i^{\text{network}}(t) = \sum_{j,f} T_{ij} \alpha(t - (t_{j,f} + \tau_{\text{ax}})) \Theta((t_{j,f} + \tau_{\text{ax}}) - t_{i,l(t)}). \quad (5)$$

The factor $\Theta((t_{j,f} + \tau_{\text{ax}}) - t_{i,l(t)})$ assures that currents due to previous synaptic inputs are reset to zero when neuron i spikes, $t_{i,l(t)}$ is the last firing time of neuron i before t , and $\Theta(x)$ is the theta function, i.e., $\Theta(x) = 0$ for $x < 0$ and $\Theta(x) = 1$ for $x \geq 0$. The term τ_{ax} denotes the axonal delay from neuron j to neuron i . Unless stated otherwise, the kernels $\alpha(\tau)$ are normalized according to

$$\int_0^\infty \alpha(\tau) d\tau = 1 \quad (6)$$

so that the terms T_{ij} are equivalent to the total integrated synaptic strength from neuron j to neuron i .

In the following two sections we study the associative capabilities of specific realizations of this general model. First we consider systems with no oscillations of the background activity and with postsynaptic currents modeled as δ -pulses. We then turn to systems with periodic background oscillations and postsynaptic currents modeled as rectangular pulses or differences of two exponential functions.

III. MODEL WITH CONSTANT BACKGROUND ACTIVITY AND FAST SYNAPTIC CURRENTS

To capture the effects of randomly impinging synaptic inputs known to put cortical neurons close to firing threshold under *in vivo* conditions [31], all model neurons receive a constant positive background current $I_i^{\text{external}}(t) = I_B > 0$ which is slightly less than one so that without recurrent or further external input, the membrane potentials relax to a level below but close to the firing threshold, $\lim_{t \rightarrow \infty} |u_{\text{thresh}} - u_i(t)| = |1 - I_B| \ll 1$. Postsynaptic currents are approxi-

mated by δ -pulses which is justified if synaptic time constants are short compared to the membrane time constant.

The kernels α are thus given by

$$\alpha(\tau) = \delta(\tau). \quad (7)$$

Let us now investigate how such a system responds to an external stimulus pattern presented to the network at time $t_0=0$. We assume input patterns that raise the membrane potentials of some neurons above firing threshold whereas the membrane potentials of all other neurons remain at their previous stationary value I_B . The set of neurons with positive input is denoted by G_0 .

The time evolution of the network can be readily understood by the following step-by-step consideration. All neurons within the group G_0 fire action potentials at time $t_0=0$. Immediately afterwards, their membrane potentials are reset to $u_i=0$. At time $t=\tau_{ax}$, the action potentials arrive at the postsynaptic neurons. Due to the constant background current $I_i^{external}(t)=I_B$ in Eq. (2), the membrane potentials of all neurons in group G_0 have reached the value $u(\tau_{ax}^-) = I_B[1 - \exp(-\tau_{ax})]$ at that time. Neurons that did not fire at time $t_0=0$ still hover at the fixed point $u=I_B$.

If the axonal delay τ_{ax} is large compared to one (the membrane time constant), all u_i are thus again just below threshold. The arrival of the action potentials fired by neurons $j \in G_0$ at time t_0 then triggers postsynaptic currents which instantaneously change the membrane potential of neuron i by $\sum_{j \in G_0} T_{ij}$. For sufficiently large τ_{ax} , neuron i will therefore fire at time τ_a if $\sum_{j \in G_0} T_{ij} > 0$ and will stay quiescent if the sum is negative. Repeating this argument shows that at the times $t_k = k\tau_{ax}$, where $k \in \mathbb{N}$, certain sets G_k of neurons are active. If we denote a neuron i that is active by $A_i=1$ and a neuron that stays quiescent by $A_i=0$, we obtain the coupled-map dynamics

$$A_i(t_{k+1}) = \Theta \left(\sum_j T_{ij} A_j(t_k) \right) \quad \text{with } t_k = k\tau_{ax}. \quad (8)$$

This time evolution is identical to the update rule of the Little model [7] with a 0/1-representation. The result implies that the group of neurons that fires action potentials at time $k\tau_{ax}$ in the integrate-and-fire model is the *same* group that is in the on-state ($A_i=1$) in the k th iteration of a Little model [7] with identical initial conditions (same group G_0 of neurons with $A_i=1$ at time t_0) and identical couplings T_{ij} .

It follows that both networks retrieve the same pattern and that eventually the group of neurons which fire remains unchanged. The duration of the transient phase depends on the number of patterns stored in the network and whether the dynamics relaxes to one of the stored patterns or to a spurious attractor [10–12]. Note that the Little model literally reaches a fixed point, i.e., the activity A_i of cell i approaches a constant value, one or zero, whereas in the present network the binary pattern is represented by a τ_{ax} -periodic firing pattern.

In the above argument, a large axonal delay τ_{ax} is required for two different reasons. First, $\tau_{ax} \gg 1$ guarantees that a drop of u_i due to a negative total recurrent input at time $(k-1)\tau_{ax}$ has decayed sufficiently by time $k\tau_{ax}$, so that u_i is again just below threshold. Second, a large delay is also re-

quired to ensure that the constant input current I_B has raised the membrane potential sufficiently close to threshold by time $k\tau_{ax}$ for a neuron which fired and was reset to $u_i=0$ at time $(k-1)\tau_{ax}$. Note that the second requirement can be avoided for arbitrary τ_{ax} by simply adding positive self-couplings $T_{ii}=I_B \exp(-\tau_{ax})$. This has been done in the numerical simulations.

It is known from analytical studies and numerical simulations that one can store an extensive number ($p \propto N$) of random patterns ξ^μ in the Little model [10–12]. The critical storage level $p_c = \alpha_c N$ is reached for $\alpha \approx 0.14$ [32]. To be precise, this result holds for a Little (or Hopfield) model with ± 1 -representation (“on” = +1, “off” = -1) whose dynamics is given by

$$S_i(t_{k+1}) = \text{sgn} \left(\sum_j T_{ij} S_j(t_k) \right). \quad (9)$$

In the Little model, all neurons are updated in parallel, in the Hopfield model [1], Eq. (9) is applied to only one neuron at a time, chosen in a serial or random sequential order.

In both models, the synaptic weights T_{ij} are determined by the Hebb rule,

$$T_{ij} = N^{-1} \sum_{\mu=1}^p \xi_i^\mu \xi_j^\mu \quad \text{for } i \neq j, \quad T_{ii} = 0, \quad (10)$$

the pattern components ξ_i^μ are chosen independently and with equal probability from $\{-1, 1\}$, i.e., $\text{Prob}(\xi_i^\mu = 1) = \text{Prob}(\xi_i^\mu = -1) = \frac{1}{2}$.

The critical storage level α_c is defined by the condition that the system remains near a stored pattern if it is initialized with that pattern, i.e., the overlap $m^\mu(t)$,

$$m^\mu(t) = N^{-1} \sum_j S_j(t) \xi_j^\mu, \quad (11)$$

remains at a value that is close to one for this particular pattern. At the critical storage level, $m_c^\mu \approx 0.97$ so that about 1.5% of the bits of a pattern are flipped, i.e., not retrieved correctly [10–12]. Above α_c , the system tends to relax to states that are only vaguely reminiscent of the initial pattern, $m^\mu \approx 0.35$. Strictly speaking, the phase transition at α_c occurs only in the limit $N \rightarrow \infty$, however, the qualitative change between the behavior below α_c and above α_c can also be seen in finite systems and may be used to determine the storage capacity in numerical studies: Below α_c , the probability to remain near a stored pattern increases with system size, above α_c it decreases [33]. As in previous studies with systems with two-state or graded-response neurons, we will use this phenomenon to determine the storage capacity of networks with integrate-and-fire neurons.

Integrate-and-fire neurons that do not fire do not have any influence on the state of the other model neurons. This implied the 0/1-representation in Eq. (8). However, most analytical results on the Little model have been obtained for the ± 1 -representation Eq. (9). In order to compare the collective dynamics of the present model with those results, we therefore transform the ± 1 -representation into the 0/1 representation by a change of variables,

$$A_i = \frac{1}{2}(S_i + 1), \quad S_i = 2A_i - 1. \quad (12)$$

Inserting Eq. (12) into Eq. (9) demonstrates that the time evolution (9) is identical to the modified dynamics

$$A_i(t_{k+1}) = \Theta \left(\sum_j T_{ij} A_j(t_k) - \frac{1}{2} \sum_j T_{ij} \right). \quad (13)$$

For patterns ξ^μ where the number of +1's balances exactly the number of -1's, the term $\sum_j T_{ij}$ vanishes if the Hebb rule (10) is used to determine the T_{ij} . For random patterns, this is only true on average. Fluctuations from the mean do, however, decrease the storage capacity of the model with the original dynamics (8).

On the level of the macroscopic iteration equations there are thus two major differences between the models; first, the occurrence of the terms $\sum_j T_{ij}$ in Eq. (13) as opposed to Eq. (8) and second, the idealization of long axonal delays implicit in Eq. (8). We are interested in the effect of realistic axonal delays on the dynamics and not on the side effect of the terms $\sum_j T_{ij}$ arising from a change of the representations used in the Little model. We therefore separate the two effects by artificially balancing the second term in Eq. (13) through an additional auxiliary neuron $i=0$. This cell is triggered by the firing activity in the network that is received after an axonal time delay τ_{ax} so that the cell also spikes at the times t_k . The synaptic strengths from this neuron to all other cells are $T_{i0} = -\frac{1}{2} \sum_j T_{ij}$ to provide the balance term in Eq. (13).

For the numerical study of the retrieval quality, $p = \alpha N$ unbiased ± 1 random patterns ξ^μ were stored using the Hebb rule Eq. (10). For given storage level α , averages from multiple realizations of networks with up to 2000 neurons were analyzed. Each simulation consisted of p runs in which the dynamics was started in states that consecutively resembled each of the stored memories. Initial overlaps $m^\mu(0)$ less than one were used to test the capability of the network to retrieve a stored pattern from incomplete or noisy inputs. Upon reaching an attractor state, the overlap of the final state with the corresponding memory was determined. In Figs. 1, 2, and 5, these overlaps are represented by histograms that display the fraction of final states with given overlap with the original pattern, averaged over all patterns and realizations. For the reader's convenience, the characteristic shape of the postsynaptic potential (PSP) in each simulation is shown as an inset in the figures, together with a sketch of the background activity $I^{\text{external}}(t)$.

In Fig. 1 the behavior of the network is compared with that of a Little model. The same storage level $\alpha = 0.145$ is chosen for both systems. In the Little model the bin with the highest overlap ($0.95 < m^\mu \leq 1.00$) slightly decreases with increasing system size and a peak at $m^\mu \approx 0.35$ develops, indicating that $\alpha = 0.145$ is already above the storage capacity α_c of the Little model, in accordance with the literature. In the present model the bin with highest overlap stays approximately constant and no peak at $m^\mu \approx 0.3$ develops, indicating that $\alpha_c \approx 0.145$. The storage capacity of the integrate-and-fire model is thus slightly higher than that of the Little model. A similar conclusion can be drawn from simulations with ini-

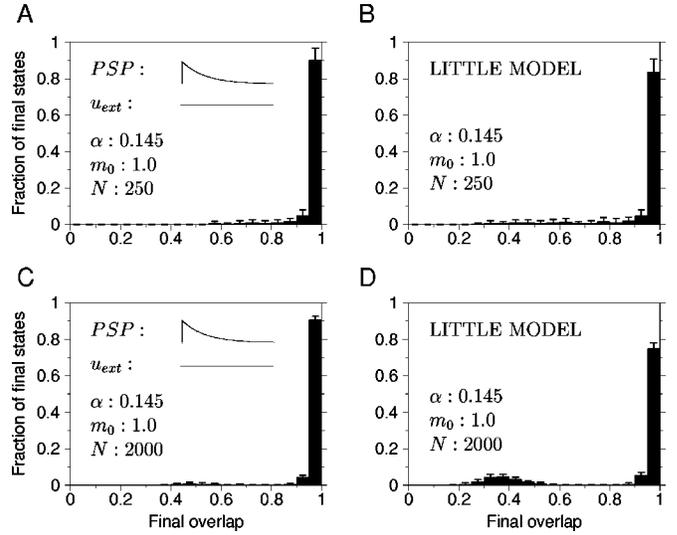


FIG. 1. Associative capability of the integrate-and-fire network with constant background activity and fast synaptic currents (left) and the Little model (right). Shown are numerical results where the network dynamics was initialized with one of the stored memory patterns and then simulated until it reached a stationary state. The fraction of final states thus obtained is plotted as a function of the overlap with the initial memory pattern (in five-percentile bins). Data are averaged over all stored patterns and 20 realizations of the synaptic coupling matrix, error bars denote standard deviations. The networks consisted of 250 neurons (A), (B) or 2000 neurons (C), (D). Comparison of (A) and (C) shows distributions of final states that are almost independent of system size, indicating that the storage capacity of the network with spiking neurons is close to $\alpha = 0.145$, the storage level used in the simulations. In (B), the fraction of final states in the highest bin is lower than in (A) and decreases with increasing system size (D). The growing peak at $m^\mu \approx 0.35$ for the Little model indicates that its storage capacity is lower than 0.145, in accordance with the literature.

tial overlaps less than one corresponding to noisy input patterns. As an example, results for $\alpha = 0.135$ and initial overlap $m_0 = 0.7$ are shown in Fig. 2.

A systematic comparison of various simulations is presented in Fig. 3. Here the fraction of final configurations with an overlap larger than 0.9 is plotted as a function of the initial overlaps. The simulations were performed for both models and storage levels below and above the storage capacity of the Little model. The results demonstrate that for low storage levels, both models exhibit very similar collective behavior. With increasing storage level, however, the network with spiking neurons performs increasingly better than the Little model. Apart from this difference, the overall similarity between the associative capabilities of both models is highly surprising if one takes into account that for the simulations, the axonal delay was chosen to be $\tau_{ax} = 0.2$, instead of $\tau_{ax} \gg 1$ required for the heuristic argument leading to Eq. (8).

Let us now try to understand the differences between both models in detail. For the chosen value of the axonal delay, $\tau_{ax} = 0.2$, the total negative recurrent input for neurons which did not fire at time $(k-1)\tau_{ax}$ has only decayed by 18.1% at time $k\tau_{ax}$, far from the assumption of a 100% decrease in Eq. (8). Therefore higher positive recurrent input for such a neuron is needed to generate an action potential at time $k\tau_{ax}$.

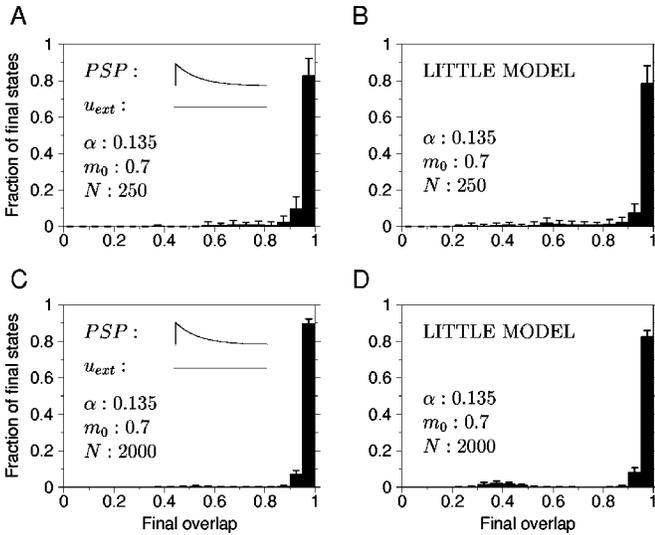


FIG. 2. Pattern completion in the integrate-and-fire network with constant background activity and fast synaptic currents (left) and the Little model (right). The networks were initialized with input patterns that were generated by randomly flipping 15% of the bits of each pattern, resulting in an initial overlap $m_0 = 0.7$. Network dynamics, numerical analysis, and representation of the data are otherwise as in Fig. 1. As shown by the simulations, run at a storage level of $\alpha = 0.135$, the network with spiking neurons has a higher noise tolerance than the Little model.

This implies that neurons which do not fire in one iteration have an increased tendency not to fire in the next cycle.

This effect is an advantage if a neuron is supposed to be quiescent in the desired memory pattern because it admits larger noise levels induced by the presence of other memories and thus allows for a higher storage capacity. However, by the same token it is a disadvantage if the neuron is supposed to be in the on-state. The simulation results reveal that the advantage is larger than the disadvantage. This phenomenon can be understood by considering that close to a stored pattern, most neurons “benefit” from the effect mentioned above. For example, let us assume that the overlap with a certain stored pattern is $m = 0.6$. In that case, about 80% of the off-neurons are supposed to be off in the memory pattern (and thus have an advantage) but only about 20% of the off-neurons are supposed to be on (and thus have a disadvantage). Furthermore, if patterns have been learned using the Hebb rule (10), the expected recurrent input to an off-neuron in a stored pattern is negative whereas the expected recurrent input to an on-neuron is positive. This implies that in a noisy pattern where some neurons are flipped the recurrent input changes: flipped on-neurons received a negative input but the absolute size of this negative input is generally less than the absolute size of the negative input an off-neuron received that is not flipped. Therefore the effect of the advantage is again larger than the effect of the disadvantage.

A number of phenomena are intimately connected with this difference between the two dynamics. First, it is well known that in the Little model trajectories that start near a stored pattern but end up in the peak at $m^\mu \approx 0.35$ show an increasing overlap with the desired target pattern during the first iterations of the dynamics, before the overlap decreases due to cross-talk from the other stored patterns [34]. Because

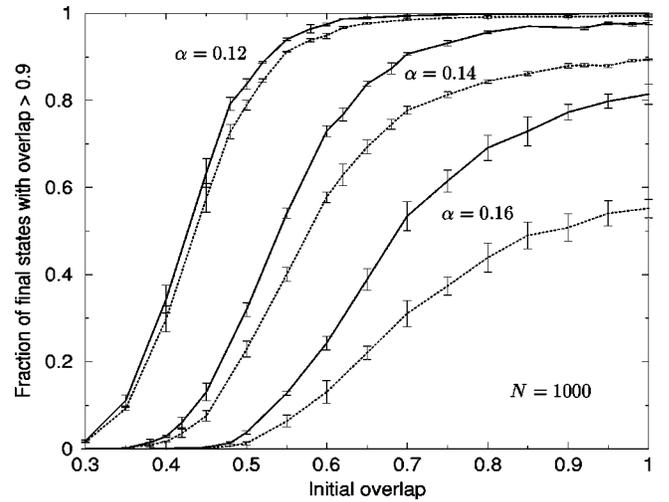


FIG. 3. Systematic comparison of pattern completion in the integrate-and-fire network with constant background activity and fast synaptic currents (solid line) and the Little model (dashed line). Both systems consisted of $N = 1000$ neurons. As a function of the initial overlap m_0 , the fraction of final configurations with an overlap $m_\infty > 0.9$, averaged over all stored pattern and 20 realizations of the coupling matrix, is plotted for three storage levels, $\alpha = 0.12$ (left curves), $\alpha = 0.14$ (middle curves), and $\alpha = 0.16$ (right curves). Vertical bars denote standard deviations. The data show that for low storage levels, the spiking and nonspiking model exhibit almost identical cooperative behavior. With increasing storage level and increasing initial overlap, the integrate-and-fire network performs increasingly better than the Little model.

of the effect mentioned above this phenomenon is less pronounced in the present model. Second, the average number of iterations needed to reach a fixed point is smaller than in the Little model, in particular for slowly relaxing solutions. Third, the difference between the performance of the two models increases with decreasing initial overlap and increasing storage level (see Fig. 3).

As argued above, the increased performance of the present model is mainly due to the improved associative behavior of neurons that are supposed to be in the off-state. The latter can be observed directly by comparing the fractions of correctly retrieved off- and on-neurons. For the Little model, the size of both fractions should be identical (apart from statistical fluctuations) and be equal to the final overlap. Plotted against the final overlap, one would thus expect that data points scatter around the diagonal but without major differences between the off- and on-neurons. Figure 4 shows that, as predicted, this is not the case for the present model: the fraction of correctly retrieved off-neurons is significantly higher than that of on-neurons.

IV. MODEL WITH SUBTHRESHOLD OSCILLATION AND LONG-LASTING POSTSYNAPTIC CURRENTS

Subthreshold membrane oscillations have been observed in many brain regions and are prominent in the hippocampus [35–38]. Oscillation frequencies usually range from a few Hertz (alpha waves) to 40–60 Hz (gamma waves). Experimental results suggest that under in vivo conditions certain classes of neurons fire at most a few action potentials in one cycle [39]. The generation of action potentials occurs mainly

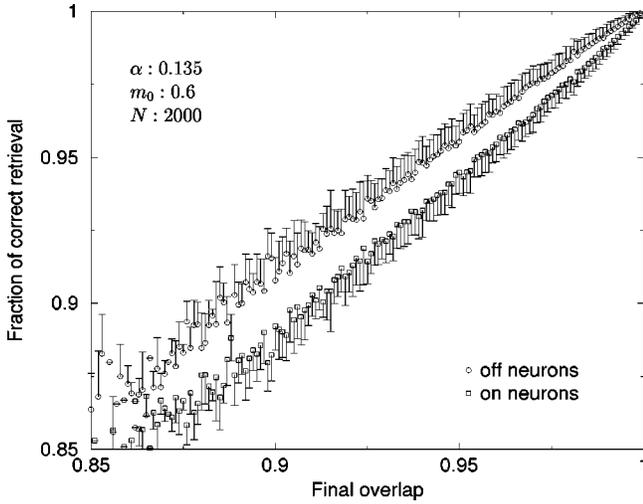


FIG. 4. Difference in the behavior of firing and quiescent neurons in the integrate-and-fire network with constant background activity and fast synaptic currents ($N=2000$, $\alpha=0.135$, $m_0=0.6$). Shown are the fractions of firing neurons (open squares) and quiescent neurons (open circles) in the final state as a function of the overlap of the final state with the stored memory pattern, averaged over all stored pattern and 20 realizations of the coupling matrix. The length of the vertical bars denotes one standard deviation. In systems without symmetry breaking between on- and off-neurons, such as the Little model, both fractions would be identical (apart from statistical fluctuations) and be equal to the final overlap. The systematic deviations demonstrate the different dynamic roles of firing and nonfiring cells in models with spiking neurons.

in the rising part of the oscillation [40]. Furthermore, depending on the specific neurotransmitter, the duration of postsynaptic potentials may be of the same order as the oscillation period. This suggests that spike activity in one oscillation cycle may generate postsynaptic potentials that last long enough to trigger action potentials far into the next cycle.

Within this paper we do not attempt to model a specific brain region or experimental paradigm. We therefore only incorporate the essence of these experimental findings into the model (2)–(6) and investigate whether the resulting system is again capable of functioning as a content-addressable memory. To account for global subthreshold oscillations of the membrane potential in a minimal way, the time-dependent part of the external input current $I^{\text{external}}(t)$ in Eq. (3) is taken to be a sine wave,

$$I^{\text{external}}(t) = I_{\text{dc}} + I_{\text{ac}} \cos(\omega t), \quad (14)$$

with period $P = 2\pi/\omega$.

As in the preceding section, model parameters are chosen such that without input from other cells or external stimuli, a neuron does not fire. The time evolution below the firing threshold is described by Eq. (2), a linear differential equation. One can therefore readily determine conditions that guarantee subthreshold behavior.

Without feedback, $I_i(t)$ equals $I^{\text{external}}(t)$, and for $t \rightarrow \infty$ all neurons approach the synchronous oscillation

$$u^{\text{osc}}(t) = I_{\text{dc}} + A \cos(\omega t + \phi), \quad (15)$$

where

$$A = A(I_{\text{ac}}, \omega) = I_{\text{ac}} \frac{1}{\sqrt{(1 + \omega^2)}} \quad (16)$$

and

$$\phi = \phi(\omega) = \arctan(\omega). \quad (17)$$

For given amplitude I_{ac} and frequency ω of the oscillating current, we thus choose the dc part I_{dc} of the input current to be slightly less than $1 - A(I_{\text{ac}}, \omega)$, i.e., $I_{\text{dc}} > 0$ and $|1 - A(I_{\text{ac}}, \omega) - I_{\text{dc}}| \ll 1$. Without loss of generality, we shift the time axis and set $\phi = 0$ so that the maxima of $u_{\text{osc}}(t)$ are reached at times kP with $k \in \mathbb{N}$.

Two different model variants will be investigated to study the dynamical effects of long-lasting postsynaptic currents. In the first variant, current pulses are approximated by rectangular pulses,

$$\alpha(\tau) = \Theta(\tau) \Theta(\tau_{\text{psc}} - \tau). \quad (18)$$

This choice implies that the change of u_i due to a single action potential generated by neuron j at time t is maximal in amplitude at time $t + \tau_{\text{ax}} + \tau_{\text{psc}}$. Its size Δu_{ij} at that time, as derived from Eq. (2), is

$$\Delta u_{ij} = [1 - \exp(-\tau_{\text{psc}})] T_{ij}. \quad (19)$$

The simple shape of the current pulse admits an analytic treatment of the retrieval behavior of the network as will be shown below. However, this choice might be oversimplified. For the second model variant, we therefore follow the literature and use the difference of two exponentials as an effective description for the time course of the postsynaptic current [41],

$$\alpha(\tau) = \alpha(\tau) = c \left[\exp\left(-\frac{\tau}{\tau_d}\right) - \exp\left(-\frac{\tau}{\tau_r}\right) \right] \Theta(\tau). \quad (20)$$

The time constant τ_r is the rise time and τ_d is the decay time of the postsynaptic current ($\tau_d \gg \tau_r$). The constant c is chosen such that the maximum of $\alpha(\tau)$ equals one as in Eq. (18). Both normalizations differ from Eq. (6) used in the preceding section and are more convenient in the present context since they allow a direct comparison of systems with different shapes and time constants of postsynaptic currents.

To mimic the experimental observation that certain classes of neurons mainly fire during the rising phase of the membrane oscillations, we consider parameter regimes such that the model neurons only fire within an interval of length D_{fire} before the maximum of the oscillation, i.e., cannot reach the threshold in the intervals $(kP, (k+1)P - D_{\text{fire}})$. Because $u^{\text{osc}}(t) \leq u^{\text{osc}}(D_{\text{fire}})$ for $t \in (kP + D_{\text{fire}}, (k+1)P - D_{\text{fire}})$, the constraint can be satisfied for arbitrary $0 \leq I_{\text{dc}} < 1 - A$ in the intervals $(kP + D_{\text{fire}}, (k+1)P - D_{\text{fire}})$ if even in the worst case where some neuron receives maximal positive input (19) the firing threshold is not reached at time $kP - D_{\text{fire}}$,

$$\max_i \left(\sum_j |T_{ij}| \right) \leq \frac{A}{1 - \exp(-\tau_{\text{psc}})} [1 - \cos(\omega D_{\text{fire}})]. \quad (21)$$

This condition can be met without loss of generality by rescaling the synaptic weights, $T_{ij} \rightarrow \gamma T_{ij}$, with an appropriate scale factor γ : this normalization has no effect on the time evolution of the Little model with which we want to compare the dynamics of the present system.

To guarantee that neurons do not fire in the remaining intervals ($kP, kP + D_{\text{fire}}$], two conditions have to be fulfilled. First, postsynaptic potentials due to action potentials generated in earlier cycles should reach their maximum and trigger a spike before the oscillation maximum is reached at time kP . This can be achieved by properly adjusting the time constant of the postsynaptic current and the membrane time constant.

Second, postsynaptic currents due to action potentials generated in a given cycle should not trigger a neuron to fire after the oscillation maximum. The sufficient (though non-optimal) condition

$$\tau_{\text{ax}} = 2D_{\text{fire}} \quad (22)$$

will be used in some of the simulations. This constraint also guarantees that neurons do not fire twice within a cycle, idealizing the experimental finding that at most a few spikes are generated in one cycle [39]. Note that Eq. (22) can be easily relaxed because long rise times of the postsynaptic potential eliminate the possibility of multiple firings within one cycle or firing after the maximum of the cycle. As in the previous model we balance the second term in Eq. (13) by adding an auxiliary spiking model neuron $i=0$ to the system. This cell does not receive input from the other neurons. By choosing an increased input current I_{dc} for this neuron it will reach threshold and fire at or slightly before the maxima of the oscillation. The synaptic strengths from this cell to all other neurons are again $T_{i0} = -1/2 \sum_j T_{ij}$.

Let us now analyze the dynamics of networks with block-shaped postsynaptic currents (18). The systems are initialized with input currents such that the membrane potentials of a group G_0 of neurons are raised above u_{thresh} at or slightly before the maximum of an oscillation cycle, whereas the other neurons receive no external input.

Once the initial spikes are generated, the membrane potentials of all neurons in G_0 are reset to zero. At time τ_{ax} , i.e., after one axonal delay time, the action potentials arrive at the postsynaptic neurons and trigger rectangular current pulses with a duration τ_{psc} that is chosen such that the postsynaptic current lasts till the maximum of the next oscillation cycle.

By the time D_{fire} before the maximum of the next oscillation cycle all the postsynaptic currents are sufficiently integrated, i.e., $\Delta u_{ij} = T_{ij} - \epsilon$ with arbitrarily small $\epsilon > 0$ if P is long enough. Thus in the limit of $P \gg 1$, a postsynaptic neuron i will fire if $\sum_{j \in G_0} T_{ij} > 0$ and will stay quiescent if the sum is negative.

Repeating this argument, we see that the group G_k of neurons that generate action potentials in the time interval D_{fire} before the maximum of the k th oscillation cycle is again identical with the group of neurons that are in the on-state in the k th iteration of a Little model [7] with the same couplings T_{ij} and initial activity pattern. This implies that as in

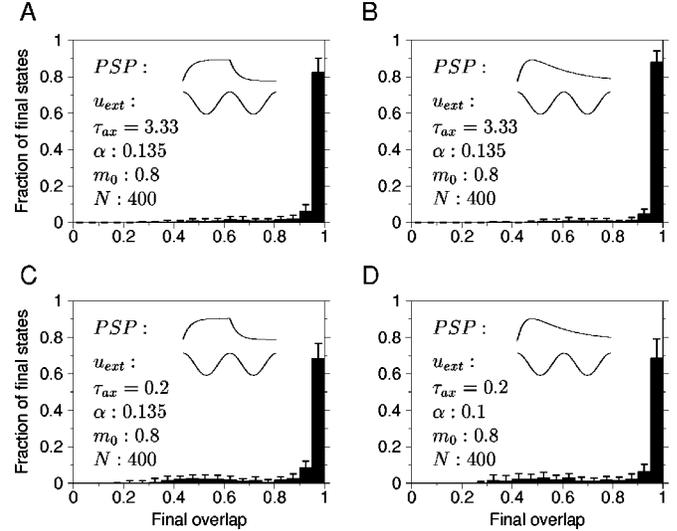


FIG. 5. Associative capability of the integrate-and-fire network with subthreshold oscillation and long-lasting postsynaptic currents. The simulations were performed in systems with $N=400$ neurons at a storage level of $\alpha=0.135$ (A), (B), (C) and $\alpha=0.1$ (D). The initial overlap with stored patterns was $m_0=0.8$. Shown are results from systems with block-shaped (A), (C) and double-exponential (B), (D) postsynaptic currents, resulting in time courses for the postsynaptic potentials as depicted in the upper insets. Axonal delays τ_{ax} were either long (A), (B) or short (C), (D) compared with the RC -time constant which was normalized to one. The oscillation period P was 6.67 and the amplitude of the oscillation was 0.4. Otherwise, network dynamics, numerical analysis, and data presentation are as in Fig. 1. Comparison of (A) with (C) or (B) with (D) demonstrates that the associative capabilities decrease with decreasing axonal delay. This effect is less pronounced for block-shaped postsynaptic currents than for postsynaptic currents with a double exponential shape — in (A) and (C), the same value for α has been used, whereas in (D) it was lowered to $\alpha=0.1$ to obtain a distribution similar to that in (C).

the preceding section, all results obtained for the Little model regarding fixed points, convergence, and storage capacity apply to the present model as well.

Note that a firing pattern in the present model consists of all neurons i that generate action potentials ($S_i=1$ in the Little model) within a time interval of fixed length D_{fire} before the maximum of one cycle of the background oscillations. In general, these neurons will not fire at the same time because of variations in the recurrent input. This situation differs from the model discussed in the preceding section where all active neurons fire in strict synchrony.

The equivalence of the present model with the Little model has now been proven for an ideal situation required for the mathematical analysis. To analyze the behavior of the network under more realistic conditions, numerical simulations with both block-shaped (18) and double-exponential postsynaptic currents (20) with physiologically realistic parameters were performed. With a membrane time constant of $\tau_{\text{RC}}=15$ ms, the values used correspond to $P=100$ ms, $\tau_d=80$ ms, $\tau_r=1$ ms, and $\tau_{\text{ax}}=3$ ms.

Some of the results are shown in Fig. 5 and can be summarized as follows. For periods P that are long compared to one (the RC -time constant) and for long axonal delays (as required for the analytic convergence proof), systems with

rectangular or double-exponential postsynaptic currents do indeed have associative capabilities similar to the Little model, as predicted by the theory.

If the axonal delay is reduced to a realistic value of $\tau_{ax} = 0.2$ (i.e., 3 ms), the storage capacity of the model decreases because neurons that fire more than τ_{ax} before the oscillation maximum may influence the firing behavior of other neurons in the same cycle. The decrease of the storage capacity is larger for the model with double-exponential current shape because the rise time to the maximum of the postsynaptic potential is shorter than for the model with rectangular postsynaptic current shape. Thus the above mentioned change in the membrane potential for the other neurons is larger. The size of the postsynaptic potential is also more strongly influenced by the precise arrival time of the presynaptic spike than in the model with rectangular current shape. This is not desirable because the system shows best pattern retrieval if the postsynaptic potentials have the same size at the time of firing.

V. DISCUSSION

The present study demonstrates that model systems with integrate-and-fire neurons can be programmed to relax for appropriate initial conditions to periodic oscillations representing stored memories. Two different mechanisms were investigated in detail. In the first case, delayed feedback within the network enforces neurons to generate action potentials at integer multiples of the feedback delay. In the second case, periodic solutions are obtained through an interplay of subthreshold oscillations and long-lasting postsynaptic potentials.

In both cases, the networks exhibit associative properties similar to those of the Little model. The results demonstrate that although the microscopic time evolution of integrate-and-fire neurons strongly differs from that of graded-response or binary neurons, the collective network dynamics may nevertheless be almost identical. On the level of macroscopic order parameters and for stationary solutions, similar findings have been reported in the literature for systems where the number p of stored patterns scales logarithmically with the number N of neurons [19]. The current study reveals that even at high storage levels ($p \propto N$) and during transient relaxations, spiking and nonspiking systems can exhibit strikingly similar cooperative phenomena.

Associative memory storage in a dynamical system requires the existence of multiple attractors. In the models studied in this paper, attractors are imprinted using a Hebbian learning rule in large networks. The attractors exhibit the same simple temporal characteristics — periodic firing patterns — but differ in their spatial structure. However, multistability can also arise in small neural systems whose dynamics allows multiple attractors that differ in their temporal signature [42]. The two scenarios highlight two options for generating multistability in neural networks: spatial versus temporal complexity. The existence of multiple concurrent rhythms in various brain structures [37] indicates that in nature both mechanisms may operate in parallel. It would therefore be interesting to link both approaches and study how complex spatiotemporal patterns can be stored at will in

extended neural systems, and how synaptic noise influences the emergent behavior of such systems.

With respect to oscillatory background activity, our findings indicate that systems with spiking neurons, realistic axonal delays, and long-lasting postsynaptic currents may utilize subthreshold oscillations of their membrane potentials to function as content-addressable memories. Thus previous approaches modeling the hippocampus as an associative memory using highly simplified networks with two-state neurons and a discrete time evolution may be justified on a phenomenological level even if they do not capture major aspects of the microscopic dynamics.

Systems with subthreshold oscillations studied in this paper demonstrate another interesting feature. When these networks settle in a periodic activity pattern, the firing time of a neuron relative to the underlying oscillation varies from neuron to neuron, and depends on the attractor reached. Cross-correlations between the spike activity of one neuron and the subthreshold oscillation or between the firings of two neurons would indicate a temporal code where information about stored patterns is encoded in the fine temporal structure of neural activity. Indeed, these temporal patterns could be used for further signal processing within the present scheme. However, these temporal phenomena do not play *any* functional role in the dynamic mechanism used for associative memory storage of the networks analyzed here—the temporal fine structure is a mere epiphenomenon of the network dynamics. This observation provides a simple example that even measuring stimulus-dependent temporal correlations in neural systems cannot be used to verify that those systems actually use the temporal domain to represent stimulus properties.

The numerical simulations showed that various idealizations required for the mathematical analysis may be violated without a significant decrease in the performance of the models as content-addressable memories. In particular, time constants governing oscillation periods may be of the same order as the membrane time constant. Let us mention in passing that for the model with constant background activity and fast synaptic currents, the collective behavior remains qualitatively the same even for much smaller values of τ_{ax} if one arranges several networks without internal recurrent connections in a staggered loop structure so that the output of one network serves as input to the subsequent network.

A study that is conceptually related to the present work shows that the dynamics of a Hopfield network [3] can be implemented in a system with spiking neurons [43]. Through a careful adjustment of the temporal characteristics of postsynaptic potentials, a neural code based on relative spike times can be accomplished. Apart from this difference in the representation of memory patterns, both studies aim at the same general goal: to realize the time evolution of one dynamical system—the Hopfield model in [43] and the Little model in the present paper—within the dynamical repertoire and constraints of a second system.

The success of the two approaches demonstrates that this goal can indeed be accomplished. However, as indicated in the Introduction, this does not imply that biological neural networks do indeed operate as content-addressable memories in the way envisaged in both studies. Similar remarks apply to theoretical studies about the dynamics and computational

capabilities of “synfire chains” [13,14,17]. Nevertheless, all these results are helpful in that they prove the principal feasibility of associative pattern storage and retrieval in systems with spiking neurons.

ACKNOWLEDGMENTS

We are grateful to David MacKay for most helpful discussions, and to Martin Stemmler for critical comments on the manuscript.

-
- [1] J.J. Hopfield, Proc. Natl. Acad. Sci. USA **79**, 2554 (1982).
 [2] M.A. Cohen and S. Grossberg, IEEE Trans. Syst. Man Cybern. **13**, 815 (1983).
 [3] J.J. Hopfield, Proc. Natl. Acad. Sci. USA **81**, 3088 (1984).
 [4] D. Marr, Proc. R. Soc. London, Ser. B **176**, 161 (1970).
 [5] D. Marr, Philos. Trans. R. Soc. London, Ser. B **262**, 24 (1971).
 [6] A. Treves and E.T. Rolls, Hippocampus **4**, 374 (1994).
 [7] W.A. Little, Math. Biosci. **19**, 101 (1974).
 [8] C.M. Marcus and R.M. Westervelt, Phys. Rev. A **40**, 501 (1989).
 [9] A.V.M. Herz and C.M. Marcus, Phys. Rev. E **47**, 2155 (1993).
 [10] D.J. Amit, *Modeling Brain Function: The World of Attractor Neural Networks* (Cambridge University Press, Cambridge, England, 1989).
 [11] J.L. van Hemmen and R. Kühn, in *Physics of Neural Networks*, edited by E. Domany, J.L. van Hemmen, and K. Schulten (Springer, Berlin, 1991), pp. 1–106.
 [12] J. Hertz, A. Krogh, and R.G. Palmer, *Introduction to the Theory of Neural Computation* (Addison-Wesley, Redwood City, 1991).
 [13] M. Abeles, *Corticonics: Neuronal Circuits of the Cerebral Cortex* (Cambridge University Press, Cambridge, England, 1991).
 [14] M. Herrmann, J.A. Hertz, and A. Prügel-Bennett, Network **6**, 403 (1995).
 [15] J.J. Hopfield, Nature (London) **376**, 33 (1995).
 [16] J.J. Hopfield and A.V.M. Herz, Proc. Natl. Acad. Sci. USA **92**, 6655 (1995).
 [17] A. Aertsen, M. Diesmann, and M.O. Gewaltig, J. Physiol. (Paris) **90**, 243 (1996).
 [18] W. Maass, Neural Comput. **9**, 279 (1997).
 [19] W. Gerstner and J.L. van Hemmen, Network **3**, 139 (1992).
 [20] R. Mueller, D.J.C. MacKay and A.V.M. Herz, in *Proceedings BioNet'96*, edited by G.K. Heinz (Gesellschaft zur Foerderung angewandter Informatik, GFaI, Berlin, 1996), pp. 70–80.
 [21] R.E. Mirollo and S.H. Strogatz, SIAM (Soc. Ind. Appl. Math.) J. Appl. Math. **50**, 1645 (1990).
 [22] Y. Kuramoto, Physica D **50**, 15 (1991).
 [23] L.F. Abbott and C. van Vreeswijk, Phys. Rev. E **48**, 1483 (1993).
 [24] M. Tsodyks, I. Mitkov, and H. Sompolinsky, Phys. Rev. Lett. **71**, 1280 (1993).
 [25] M. Usher, H. Schuster, and E. Niebur, Neural Comput. **5**, 570 (1993).
 [26] C. van Vreeswijk and L.F. Abbott, SIAM (Soc. Ind. Appl. Math.) J. Appl. Math. **53**, 253 (1993).
 [27] D. Hansel, G. Mato, and C. Meunier, Neural Comput. **7**, 307 (1995).
 [28] U. Ernst, K. Pawelzik, and T. Geisel, Phys. Rev. Lett. **74**, 1570 (1995).
 [29] A. Nischwitz and H. Glünder, Biol. Cybern. **73**, 389 (1995).
 [30] G.J. Stuart and B. Sakmann, Nature (London) **367**, 69 (1994).
 [31] W.R. Softky, Curr. Opin. Neurobiol. **5**, 239 (1995).
 [32] T. Stiefvater, K.R. Müller, and R. Kühn, Physica A **232**, 61 (1996).
 [33] D.J. Amit, H. Gutfreund, and H. Sompolinsky, Ann. Phys. (N.Y.) **173**, 30 (1987).
 [34] H. Horner, D. Bormann, M. Frick, H. Kinzelbach, and A. Schmidt, Z. Phys. B **76**, 381 (1989).
 [35] T.H. Brown and A.M. Zador, in *The Synaptic Organization of the Brain*, edited by G.M. Shepherd (Oxford University Press, New York, 1990), pp. 346–388.
 [36] S.R. Cobb, E.H. Buhl, K. Halasky, O. Paulsen, and P. Somogyi, Nature (London) **378**, 75 (1995).
 [37] G. Buzsaki and J.J. Chobrak, Curr. Opin. Neurobiol. **5**, 504 (1995).
 [38] Y. Gutfreund, Y. Yarom, and I. Segev, J. Physiol. (Paris) **483**, 621 (1995).
 [39] K.J. Jeffrey, J.G. Donnett, and J. O'Keefe, NeuroReport **6**, 2166 (1995).
 [40] S.E. Fox, S. Wolfson, and J.B. Ranck, Jr., Exp. Brain Res. **62**, 495 (1986).
 [41] W. Rall, J. Neurophysiol. **30**, 1138 (1967).
 [42] J. Foss, F. Moss, and J. Milton, Phys. Rev. E **55**, 4536 (1997).
 [43] W. Maass and T. Natschlagler, Network **8**, 355 (1997).