

Associative Memory with Networks of Spiking Neurons in Temporal Coding

Wolfgang Maass, Thomas Natschläger
Institute for Theoretical Computer Science
Technische Universität Graz
Klosterwiesgasse 32/2, A-8010 Graz, Austria
E-mail: {maass, tnatschl}@igi.tu-graz.ac.at

Abstract

A theoretical model for analog computation in networks of spiking neurons with temporal coding is introduced and tested through simulations in GENESIS. It turns out that the use of multiple synapses yields very noise robust mechanisms for analog computations via the timing of single spikes in networks of detailed compartmental neuron models. One arrives in this way at a method for emulating arbitrary Hopfield nets with spiking neurons in temporal coding, yielding new models for associative recall of spatio-temporal firing patterns.

1 Introduction

Recent experimental results from neurophysiology have shown that in many biological neural systems not only the firing rate, but also the spatio-temporal pattern of neuronal firing carries important information.

Other recent experimental results indicate that it is in fact questionable whether biological neural systems are *able* to carry out analog computation with analog variables represented as firing rates. Due to “synaptic depression”¹ the amplitude of postsynaptic potentials arising from a presynaptic neuron u tends to scale like $1/f$ where f is the firing rate of u . Therefore both slowly firing neurons and rapidly firing neurons u inject roughly the same amount of current into a postsynaptic neuron during a given time window. This suggests that both a McCulloch-Pitts neuron and a sigmoidal neuron model overestimate the computational capability of a biological neuron for rate coding.

In addition it has been argued that in view of the rather low firing rates of cortical neurons analog computations in multi-layer neural systems with intermediate variables represented as firing rates would be much too slow to

achieve the experimentally observed computation speed of concrete cortical neural systems^{2,3}.

Temporal coding with single spikes has been suggested as an alternative neural code, that would not be affected by synaptic depression. In addition with this neural code one can in principle achieve very high computation speed with biologically realistic low firing rates. In temporal coding with single spikes one can encode analog variables in the firing times of neurons relative to the stimulus onset⁴, relative to some oscillation⁵, or relative to the firing times of other neurons^{6,2}.

In this article we explore models for associative memory in networks of spiking neurons with the latter two types of temporal coding. In this regard our models differ from previous models for associative memory in networks of spiking neurons, that were based on rate coding^{7,8,9,10}. We exhibit both through a rigorous mathematical result and through simulations of compartmental neuron models in GENESIS a computational mechanism that allows us to emulate any given Hopfield net with a network of spiking neurons with temporal coding. This mechanism can be implemented in a surprisingly noise robust way both with regard to some underlying background oscillation⁵ and in a self-excitatory manner without any external “clock”.

Apart from these specific results, this article also addresses an important general question regarding the modeling of neural computation: To what extent are mechanisms and results that have been demonstrated analytically for networks of relatively simple mathematical models for leaky integrate-and-fire neurons also valid for networks of substantially more complex compartmental neuron models?

Section 2 of this article presents a theoretical result regarding the emulation of Hopfield nets in networks of spiking neurons. In section 3 we discuss the results of our GENESIS-simulations of related networks of compartmental neuron models. Section 4 contains some conclusions.

2 Theoretical Result

In this section we outline the construction of a network $\mathcal{S}_{\mathcal{H}}$ of spiking neurons which approximates the computation of an arbitrary given Hopfield network \mathcal{H} . As a model for a spiking neuron we take the common model of a *leaky integrate-and-fire neuron with noise*, respectively the somewhat more general *spike response model*^{9,11}. The only specific assumption that is needed for the construction of $\mathcal{S}_{\mathcal{H}}$ in Theorem 2.1 is that both the beginning of the rising part of an EPSP and the beginning of the descending part of an IPSP can be described by a linear function.




Figure 1: Piecewise linear activation function of the Hopfield net \mathcal{H} with graded response.

Theorem 2.1 *Let \mathcal{H} be an arbitrary given Hopfield net with graded response¹² and synchronous update. We assume that \mathcal{H} consists of n sigmoidal neurons u_i for $i \in \{1, \dots, n\}$ with arbitrary weights $w_{ij} \in \mathbb{R}$ for $i, j \in \{1, \dots, n\}$ and a piecewise linear activation function σ as indicated in Figure 1. Then one can approximate any computation of \mathcal{H} by a recurrent network $\mathcal{S}_{\mathcal{H}}$ of n spiking neurons (with $O(1)$ auxiliary spiking neurons) in temporal coding. An input, internal state, or output of \mathcal{H} of the form $\langle x_1, \dots, x_n \rangle \in [-1, 1]^n$ is represented in $\mathcal{S}_{\mathcal{H}}$ by temporal coding, i.e. by a firing pattern of $\mathcal{S}_{\mathcal{H}}$ in which its i -th neuron fires at time $kT - c\tilde{x}_i$, where $|x_i - \tilde{x}_i|$ can be made arbitrarily small. The reference time points kT for $k = 0, 1, \dots$ are defined by $O(1)$ periodically firing auxiliary neurons in $\mathcal{S}_{\mathcal{H}}$. Any fixpoint of \mathcal{H} corresponds to a stable periodic firing pattern of $\mathcal{S}_{\mathcal{H}}$.*

Proof: We will have in $\mathcal{S}_{\mathcal{H}}$ spiking neurons v_1, \dots, v_n that simulate the output of the sigmoidal neurons u_1, \dots, u_n in \mathcal{H} through the time at which they fire. We assume that there exists a synapse from neuron v_j to neuron v_i for any $i \neq j$. The spiking neurons v_1, \dots, v_n in $\mathcal{S}_{\mathcal{H}}$ fire periodically, for the k -th time during the time interval $[kT - c, kT + c]$, $k \in \mathbb{N}$. When the neurons v_1, \dots, v_n in $\mathcal{S}_{\mathcal{H}}$ fire during the k -th time interval at times $kT - cx_i(k)$, $i \in \{1, \dots, n\}$, then this temporal firing pattern will correspond to an output $\mathbf{x}(k) = \langle x_1(k), \dots, x_n(k) \rangle \in [-1, 1]^n$ of the n sigmoidal neurons of \mathcal{H} after their k -th parallel computation step. These values are updated in \mathcal{H} according to the equation

$$x_i(k+1) = \sigma \left(\sum_{j=1}^n w_{ij} x_j(k) \right), \quad (1)$$

where σ is the activation function of \mathcal{H} (see Figure 1) and w_{i1}, \dots, w_{in} are the weights of u_i in \mathcal{H} . The vector $\langle x_1(0), \dots, x_n(0) \rangle \in [-1, 1]^n$ is the network input.

Thus to simulate the $(k + 1)$ -th parallel computation step of \mathcal{H} in $\mathcal{S}_{\mathcal{H}}$, we need to make sure that each spiking neuron v_i fires at time $t_i(k + 1) = (k + 1)T - c\tilde{x}_i(k + 1)$ with $|\tilde{x}_i(k + 1) - x_i(k + 1)|$ as small as possible.

We achieve this by exploiting in $\mathcal{S}_{\mathcal{H}}$ a mechanism for computing a weighted sum in temporal coding. This simulation¹¹ is based on the observation that in the presence of some other excitation which moves the membrane potential close to the firing threshold, individual EPSP's respectively IPSP's (or volleys of synchronized PSP's) are able to *shift* the firing time of a neuron. This mechanism is particularly easy to analyze if we work in a range where all PSP's can be approximated well by linear functions. For this range one can show that the resulting firing time is linearly related to the *weighted sum* of the firing times of the presynaptic neurons – with the weights corresponding to the *efficacies* ("strengths") of the involved synapses.

To be precise, the firing time $t_i(k + 1)$ of neuron v_i during the simulation of the $(k + 1)$ -th parallel computation step of \mathcal{H} can be expressed as follows.

$$t_i(k + 1) = kT + d + \frac{\theta}{\sum_{j=1}^n \alpha_{ij}} - \frac{\sum_{j=1}^n \alpha_{ij}(kT - t_j(k))}{\sum_{j=1}^n \alpha_{ij}} \quad (2)$$

In this equation θ is the threshold of the spiking neuron v_i , d the propagation delay (time from generation of a spike at neurons v_j to the onset of the PSP at the soma at neuron v_i) and α_{ij} is the slope of the linear increasing (decreasing) section of the corresponding EPSP (IPSP) at the soma of v_i . The preceding equation (2) holds rigorously if neuron v_i fires at a time $t_i(k + 1)$ so that each postsynaptic potential from the previous round of firing of v_1, \dots, v_n is at time $t_i(k + 1)$ in its initial linearly rising phase at the trigger zone of v_i (in the case of an EPSP) or in its initial linearly declining phase (in the case of an IPSP).

The term $\sum_{j=1}^n \alpha_{ij}$ in the denominator normalizes in equation (2) the sum of the effective weights $\varrho' = \alpha_{ij} / \sum_{j=1}^n \alpha_{ij}$ in the weighted sum of the terms $kT - t_j(k)$ to 1. To simulate networks \mathcal{H} with weights that do not sum up to 1 one is forced to use in $\mathcal{S}_{\mathcal{H}}$ an additional neuron v_0 . If v_0 fires at times kT , $k \in \mathbb{N}$, one can then satisfy equation (3) for the next firing time $t_i(k + 1)$ of neuron v_i . One sets $T := d + \theta/\lambda$ and one chooses values α_{ij} for $j = 0, \dots, n$ so that $\sum_{j=0}^n \alpha_{ij} = \lambda$ and $\alpha_{ij} = \lambda w_{ij}$, where $\lambda > 0$ can be chosen arbitrarily. Since $\tilde{x}_j(k)$ is defined through the equation $kT - t_j(k) = c\tilde{x}_j(k)$, one gets

$$t_i(k + 1) = (k + 1)T - c \sum_{j=1}^n w_{ij} \tilde{x}_j(k). \quad (3)$$

Hence the firing time $t_i(k + 1) = (k + 1)T - c\tilde{x}_i(k + 1)$ of v_i encodes the value $\tilde{x}_i(k + 1) = \sum_{j=1}^n w_{ij} \tilde{x}_j(k)$ in temporal coding. In other words, ac-

ording to equation (3) neuron v_i computes in *temporal* coding a weighted sum of the values $\tilde{x}_j(k)$. The above only holds for values of $\sum_{j=1}^n w_{ij} \tilde{x}_j(k)$ in the linear range $[-1, 1]$ of the activation function σ . In order to achieve that v_i fires for any $\langle \tilde{x}_1(k), \dots, \tilde{x}_n(k) \rangle \in [-1, 1]^n$ approximately at time $(k+1)T - c\sigma\left(\sum_{j=1}^n w_{ij} \tilde{x}_j(k)\right)$, one can employ in the formal construction of $\mathcal{S}_{\mathcal{H}}$ some auxiliary neurons that prevent a firing of v_i during the intermediate time interval $(kT + c, (k+1)T - c)$ and which make sure that v_i definitely fires once during the time interval $[(k+1)T - c, (k+1)T + c]$.

In our formal construction these auxiliary neurons in $\mathcal{S}_{\mathcal{H}}$ have an undesired side effect on the firing time $t_i(k+1)$ if $\sum_{j=1}^n w_{ij} \tilde{x}_j(k) \in [-1, 1]$ and the value of this sum is close to -1 or +1. The resulting firing time $t_i(k+1)$ of v_i in $\mathcal{S}_{\mathcal{H}}$ is shifted due the PSP's generated by this auxiliary neurons. This is the reason why the computation of \mathcal{H} cannot be simulated *precisely* in $\mathcal{S}_{\mathcal{H}}$. But the amount of this shift can be made arbitrarily small. The details concerning these auxiliary neurons as well as noise and refractory behavior are described elsewhere¹¹.

3 Simulations with GENESIS

In the remainder of this article we take the preceding theoretical construction as basis for a case study. We want to find out to what extent mechanisms for computations with temporal coding that have been verified theoretically for integrate-and-fire neurons correspond to stable computational mechanisms for substantially more detailed compartmental models of biological neurons simulated in GENESIS¹³.

Surprisingly, it turns out that the most essential computational mechanism that underlies the proof of Theorem 2.1 works in the more detailed neuron model of GENESIS *even better*. The previously described computational mechanism for computing a weighted sum $\sum_{j=1}^n w_{ij} x_j(k)$ is theoretically sound as long as during each wave of firing the range $[0, 2c]$ of the differences in firing times of the n presynaptic neurons v_1, \dots, v_n is so small that there is a time point at which — in spite of their temporal differences — the resulting post-synaptic potentials are at the soma of each v_i *all* in their initial *linear* phase. However, since non-NMDA EPSP's rise very fast, the theory would suggest that the length $2c$ of this interval would have to be chosen around 1 ms. This is so small that one has to be concerned about the effect of various sources of temporal jitter on the precision of this temporal coding. Figure 2 shows that for GENESIS-neurons the value $2c$ can be chosen much larger.

One can extend this length by replacing each synapse between a presynaptic neuron v_j and v_i by $l > 1$ synapses. In this way a single spike from




Figure 2: Superposition of several non-NMDA EPSP's caused by the firing of a single presynaptic neuron with multiple synapses. The membrane voltage is measured at the dendrite (D) and at the soma (S). Three cases are shown: a single synapse (A), three synapses (B) and six synapses (C) from the presynaptic neuron. In all subsequent simulations we use three synapses, which results in a time interval of 7 ms during which the superposition of these EPSP's increases linearly.

v_j causes a superposition of l EPSP's at the soma of v_i . This superposition may have a substantially longer increasing phase if the signal pathways along the l synapses between v_j and v_i have a reasonable difference in their total delays (time from generation of a spike at neuron v_j to the onset of the PSP at the soma at neuron v_i). Beside from the differences in propagation times in the dendritic tree¹⁴ also the differences of the arrival times at various terminals of the axonal tree¹⁵ contribute to the diversity of the total delays (in our simulations we assume a difference of 3.5 ms). However the simple models for leaky integrate-and-fire neurons predict that the resulting rising phase of a superposition of several EPSP's is quite "bumpy", which limits its use for the here considered computational mechanism. In simulations with GENESIS the nonlinearity of this rising phase is clearly visible if one measures the resulting membrane voltages close to the synapses (see the curves labeled "D" in Figure 2). However at the *soma*, where the shape of this superposition becomes *relevant* for the here considered *computational mechanism*, this *rising phase* of the superposition of several EPSP's is *almost perfectly linear* (see the curves labeled "S" in Figure 2). This observation allows us in our GENESIS-simulations to stretch the length $2c$ of the interval for temporal coding through the use of multiple synapses to $2c = 4$ ms (part C of Figure 2 shows that a much larger value can be chosen if more than 3 synapses are used and a difference of the delays of 10 ms is assumed).

Now we introduce a method for constructing for any given Hopfield net \mathcal{H} a network $\mathcal{G}_{\mathcal{H}}$ of biological quite detailed spiking neurons so that $\mathcal{G}_{\mathcal{H}}$ has an

equally rich set of attractors as \mathcal{H} . Note that each attractor of $\mathcal{G}_{\mathcal{H}}$ is a spatio-temporal firing pattern, whereas an attractor (fixpoint) in \mathcal{H} has no temporal structure. Furthermore we show that these firing patterns of $\mathcal{G}_{\mathcal{H}}$ have in fact attractor qualities: if one initiates $\mathcal{G}_{\mathcal{H}}$ with a firing pattern that differs from all of the stable firing patterns it moves to an attractor firing pattern within a few firing periods.

The “blueprint” for the construction of $\mathcal{G}_{\mathcal{H}}$ is provided by the formal model $\mathcal{S}_{\mathcal{H}}$ that we have described in section 2. There are minor changes needed to move the construction of $\mathcal{S}_{\mathcal{H}}$ into a parameter range that is biologically more realistic.

One detail of our case study concerns the simulation of the *saturated* segments of the activation function σ (i.e. the needed “nonlinearity”) of the given artificial neural net \mathcal{H} . Whereas auxiliary neurons are needed for that purpose in the theoretical construction of Theorem 2.1, the natural form of EPSP’s and IPSP’s suggests that these auxiliary neurons may be unnecessary in a practical context¹¹. This conjecture is supported through all our simulations of $\mathcal{G}_{\mathcal{H}}$, where no extra mechanism forces a neuron to fire once during each firing wave.

For our GENESIS-simulations we started out from a Hopfield net \mathcal{H} with $n = 60$ neurons, whose weights were computed with the projection method¹⁶ so that 9 randomly chosen vectors from $\{-1, 1\}^{60}$ are fixpoints of \mathcal{H} . We have simulated the network $\mathcal{S}_{\mathcal{H}}$ of spiking neurons (that simulates \mathcal{H} according to Theorem 2.1) in GENESIS with a network $\mathcal{G}_{\mathcal{H}}$ of 61 neurons. In $\mathcal{G}_{\mathcal{H}}$ each neuron is modeled with 122 compartments, and there are 3 synapses between each pair of neurons (whose delays differ by up to 3.5 ms). $\mathcal{G}_{\mathcal{H}}$ has the same architecture and the same weights (properly scaled) as $\mathcal{S}_{\mathcal{H}}$ (30% of the weights were rounded to 0). Weights with negative values are modeled by inhibitory synapses.

In the protocols of our simulations of $\mathcal{G}_{\mathcal{H}}$ the bars to the left (“i-error”) indicate the difference between the firing times of the *first* wave of firing and the closest memory pattern (both interpreted as vectors from $[-1, +1]^{60}$ via temporal coding). Correspondingly the bars to the right (“o-error”) indicate the difference between the firing times of the *last* shown firing wave and the same memory pattern. In both cases a non-firing of a neuron is treated like a firing at the very end of the firing wave. This is justified by the observation that both scenarios reflect equivalent ways of encoding the smallest possible analog value “-1”¹¹. A neuron v_j that fires *very late* during a firing wave has an equally negligible impact on the firing times during the next firing wave as if v_j would not have fired at all. It is interesting to note that in all our simulation both types of encoding “-1” are present. During the first few firing waves “-1” is encoded by a late firing whereas after the network has reached




Figure 3: GENESIS-simulation of $\mathcal{G}_{\mathcal{H}}$ with a corrupted version \mathbf{x} of memory pattern 1 as input. The two lower diagrams show the firing times of the first 10 neurons at a higher temporal resolution. The input \mathbf{x} was constructed as follows from memory pattern 1: randomly chosen 15% of the components of memory pattern 1 were multiplied with -1 . On top of this, the value of each component of \mathbf{x} was moved by a random value from $[-0.4, 0.4]$. This input \mathbf{x} was presented to $\mathcal{G}_{\mathcal{H}}$ in *temporal coding*, although not precisely: randomly chosen 10% of the neurons in $\mathcal{G}_{\mathcal{H}}$ were prevented from firing (corresponding to an input value -1). The resulting effective difference between the *input* vector $\mathbf{x} \in [-1, +1]^{60}$ that is given to $\mathcal{G}_{\mathcal{H}}$ in temporal coding and memory pattern 1 is indicated to the left (“i-error”). For each neuron i an error bar of length in $[0, 1]$ indicates the size of the deviation between the i -th component of \mathbf{x} and the i -th component of memory pattern 1. The difference between the *output* vector $\mathbf{y} \in [-1, +1]^{60}$ in the same temporal coding (with non-firing viewed as an encoding of -1) and memory pattern 1 is indicated to the right (“o-error”).




Figure 4: Membrane potential at the soma of neuron 5 in the simulation from Figure 3 with and without noise. The diagram *with* noise reflects the additional internal noise underlying *all* our GENESIS-simulations. We assume the existence of a random current at the soma of each neuron (Gaussian distribution with a variance of 10^{-18} A^2), and that each synapse fails with a probability of 15%.

an attracting firing pattern with very high preference the neurons choose to encode a negative value close to -1 by *not* firing at all. Thus a subsequent network is able to read out the result of the computation of $\mathcal{G}_{\mathcal{H}}$ in the more noise robust form of a *rate* code.

In the space right before “o-error” we have marked in Figures 3 and 5 with a horizontal bar each neuron i of $\mathcal{G}_{\mathcal{H}}$ which does *not* satisfy following property: Neuron i fires during the last firing wave *if and only if* the i -th component of the memory pattern is a “+1”. One sees that only relatively low “errors” occur where a neuron does not satisfy this property.

To investigate the noise robustness of our model two internal sources of noise are considered (in addition to the noise by which the network inputs are perturbed): a noisy membrane potential and failing synapses (see Figure 4). These two types of noise are underlying all our simulations, which shows that our construction is robust against a substantial level of noise.

The results of the simulation of $\mathcal{G}_{\mathcal{H}}$ with a *very* noisy version of memory pattern 1 as input is shown in Figure 3. Further diagrams (published elsewhere¹⁷) show that in fact *all arbitrarily chosen fixpoints of \mathcal{H}* can be associatively recalled by $\mathcal{G}_{\mathcal{H}}$. Figure 3 shows that after 10 firing waves the memory pattern is recalled in the “digitalized” output code (firing/non-firing) *without any error*. Furthermore even after 5 firing waves almost no “digital” errors in associative recall in terms of firing/non-firing occur.

It turns out that in contrast to the theoretical model $\mathcal{S}_{\mathcal{H}}$ no separate




Figure 5: Here the input to $\mathcal{G}_{\mathcal{H}}$ is a linear combination of memory pattern 1 with factor 0.5 and of memory pattern 4 with factor 0.3, presented in temporal coding, with noise added as in Figure 3. The error bars to the left indicate the difference between the actual input and memory pattern 1 – which is its stronger component. The error bars to the right indicate the difference between the output and memory pattern 1 in temporal coding.

“oscillator” is needed in $\mathcal{G}_{\mathcal{H}}$ to define the reference times kT for $k = 0, 1, 2, \dots$. Instead, neuron 0 of $\mathcal{G}_{\mathcal{H}}$, whose firing times provide the reference times kT , is a neuron like all the others, and it receives inputs from all other neurons in $\mathcal{G}_{\mathcal{H}}$ through synapses with equal weight. It also has outgoing synapses to all other neurons in $\mathcal{G}_{\mathcal{H}}$. Thus its firing contributes to the triggering of the next firing wave.

The network $\mathcal{G}_{\mathcal{H}}$ is not only capable of an associative recall of memory patterns if the input consists of a noisy version of one of the stored patterns. It also has the ability to find the stronger component if the stimulus is a combination of two of the stored memory patterns. This fact is demonstrated through simulations. The results are shown in Figure 5. This is in contrast to a pattern recognition system comprised of spiking neurons proposed by Hopfield. Hopfield’s construction⁵ is based on “grandmother neurons” that each encode one spatio-temporal firing pattern through the locations of synapses on their dendritic tree.

In $\mathcal{G}_{\mathcal{H}}$ we employ large transmission delays of 25 ms in order to avoid refractory effects and we modeled each negative weight w_{ij} from the network \mathcal{H} through separate inhibitory synapses. One may argue that these two assumptions are biologically rather implausible. It is shown that one can construct a

layered version of $\mathcal{G}_{\mathcal{H}}$ with both shorter delays and using inhibition mainly as a control to the total network activity¹⁷.

4 Conclusions

We have exhibited a new way of simulating Hopfield nets with graded response by recurrent networks of spiking neurons with temporal coding. The resulting networks of spiking neurons carry out their computation very fast, even with biologically realistic low firing rates.

We have first demonstrated this simulation method through a rigorous theoretical result. We have then chosen the resulting networks as basis for a case study, where we check the essential computational mechanisms of a common simplified theoretical model through simulations in GENESIS. It turns out that our construction method yields surprisingly noise-robust mechanisms for analog computations by compartmental neuron models with temporal coding. In particular the dynamic range that is available for temporal coding can be made quite large by using multiple synapses between neurons. In this way one can work in a range where the temporal coding scheme is not seriously affected by realistic values of temporal jitter in firing times.

Our simulations in GENESIS show that one can in fact simulate an arbitrarily given Hopfield net through a network of compartmental neuron models with temporal coding. The resulting network has even under the influence of rather realistic amounts of internal noise a rich set of different spatio-temporal firing patterns as attractors. These attractors mimic the set of attractors of the given Hopfield net.

Acknowledgments

This work was supported by the “Fonds zur Förderung der wissenschaftlichen Forschung (FWF), Austrian Science Fund”, project number P12153.

References

1. L. F. Abbott, J. A. Varela, K. Sen, and S. B. Nelson. Synaptic depression enhances information transfer between cortical neurons. preprint, 1996.
2. S. T. Thorpe and M. Imbert. Biological constraints on connectionist modelling. In R. Pfeifer, Z. Schreter, F. Fogelman-Soulie, and L. Steels, editors, *Connectionism in Perspective*. Elsevier, 1989.

3. D. I. Perret, E. T. Rolls, and W. C. Caan. Visual neurons responsive to faces in the monkey temporal cortex. *Experimental Brain Research*, 47:329–342, 1982.
4. T. W. Kjaer, T. J. Gawne, and B. J. Richmond. Latency: another potential code for feature binding in striate cortex. *Journal of Neurophysiology*, 76(3), 1996.
5. J. Hopfield. Pattern recognition computing using action potential timing for stimulus representation. *Nature*, 376:33–36, 1995.
6. J. O’Keefe and M. L. Reece. Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus*, 3(3):3317–30, 1993.
7. E. Fransén. *Biophysical Simulation of Cortical Associative Memory*. PhD thesis, Stockholm University, October 1996.
8. M. W. Simmen, E. T. Rolls, and A. Treves. Rapid retrieval in an autoassociative network of spiking neurons. In J. M. Bower, editor, *Computational Neuroscience*, pages 273–278. Academic Press, London New York, 1995.
9. W. Gerstner and J. L. van Hemmen. Associative memory in a network of “spiking” neurons. *Network*, 3:139–164, 1992.
10. A. V. M. Herz, Z. Li, and J. L. van Hemmen. Statistical mechanics of temporal association in neural networks with transmission delays. *Physical Review Letters*, 66(10):1370–1373, 1991.
11. W. Maass. Fast sigmoidal networks via spiking neurons. *Neural Computation*, 9:279–304, 1997.
12. J. J. Hopfield. Neurons with graded response have collective computational properties like those of two-state neurons. *Proceedings of the National Academy of Sciences, USA*, 81:3088–3092, 1984.
13. J. M. Bower and D. Beeman. *The Book of GENESIS: Exploring Realistic Neural Models with the GEneral NEural SIMulation System*. Springer-Verlag, Inc. Published by TELOS, New York, 1995.
14. A. Zador, H. Agmon-Snir, and I. Segev. The morphoelectronic transform: A graphical approach to dendritic function. *The Journal of Neuroscience*, 15(3):1669–1682, 1995.
15. Y. Manor, C. Koch, and I. Segev. Effect of geometrical irregularities on propagation delay in axonal trees. *Biophysical Journal*, 60:1424–1437, 1991.
16. J. Hertz, A. Krogh, and R. G. Palmer. *Introduction to the Theory of Neural Computation*. Addison-Wesley, 1991.
17. W. Maass and T. Natschläger. Networks of spiking neurons can emulate arbitrary Hopfield nets in temporal coding. *Network: Computation in*

Neural Systems, 8(4):355-372, 1997.