PERGAMON

Neural Networks 14 (2001) 825–834

Neural Networks

www.elsevier.com/locate/neunet

2001 Special issue

Associative memory in networks of spiking neurons

Friedrich T. Sommer[a,*], Thomas Wennekers[b]

[a]University of Ulm, Department of Neural Information Processing, D-89069 Ulm, Germany
[b]Max-Planck Institute for Mathematics in the Sciences, Inselstraße 22-26, D-04103 Leipzig, Germany

Received 10 October 2000; revised 28 March 2001; accepted 28 March 2001

Abstract

Here, we develop and investigate a computational model of a network of cortical neurons on the base of biophysically well constrained and tested two-compartmental neurons developed by Pinsky and Rinzel [Pinsky, P. F., & Rinzel, J. (1994). Intrinsic and network rhythmogenesis in a reduced Traub model for CA3 neurons. *Journal of Computational Neuroscience*, *1*, 39–60]. To study associative memory, we connect a pool of cells by a structured connectivity matrix. The connection weights are shaped by simple Hebbian coincidence learning using a set of spatially sparse patterns. We study the neuronal activity processes following an external stimulation of a stored memory. In two series of simulation experiments, we explore the effect of different classes of external input, tonic and flashed stimulation. With tonic stimulation, the addressed memory is an attractor of the network dynamics. The memory is displayed rhythmically, coded by phase-locked bursts or regular spikes. The participating neurons have rhythmic activity in the gamma-frequency range (30–80 Hz). If the input is switched from one memory to another, the network activity can follow this change within one or two gamma cycles. Unlike similar models in the literature, we studied the range of high memory capacity (in the order of 0.1 bit/synapse), comparable to optimally tuned formal associative networks. We explored the robustness of efficient retrieval varying the memory load, the excitation/inhibition parameters, and background activity. A stimulation pulse applied to the identical simulation network can push away ongoing network activity and trigger a phase-locked association event within one gamma period. Unlike as under tonic stimulation, the memories are not attractors. After one association process, the network activity moves to other states. Applying in close succession pulses addressing different memories, one can switch through the space of memory patterns. The readout speed can be increased up to the point where in every gamma cycle another pattern is displayed. With pulsed stimulation, bursts become relevant for coding, their occurrence can be used to discriminate relevant processes from background activity.
© 2001 Elsevier Science Ltd. All rights reserved.

Keywords: Compartment-neuron model; Sparse associative memory; Long-term memory; Gamma oscillations

1. Introduction

Investigating the brain from a computational point of view leads to three difficult questions that have to be answered at the same time: (1) What information is processed? (2) How is the information coded in the spatio-temporal structure of neuronal activity? (3) What computational operations are performed by cortical networks? Simulation studies with artificial computational neural network models have become a tool to assist experimental brain research by testing hypothetical structure–function relationships. Associative memory is perhaps the best studied computational function of neural networks, where many quantitative studies have been performed in models with various degrees of abstraction. Although outlining possible strategies of distributed network computation in the cortex, most of these studies employ too abstracted neuronal models to understand the fast temporal structure of cortical activity, such as synchronicity between single spikes, gamma oscillations and bursting, see also the discussion of attractor network models with varying degree of realism in Sommer and Wennekers (2000).

The questions raised can be examined by simulation experiments employing biophysically faithful compartment-neurons with a description of the active membrane properties à la Hodgkin–Huxley. There are a couple of compartmental-neuron associative memory studies in the literature. Lansner and Fransén (1992) investigated the retrieval of learned patterns by unsynchronized rate coding persistent over several hundreds of milliseconds. In a later work, the same authors viewed not the single cortical cells, but small cell groups within columns as the rate coding subunits of the cortex (Fransén & Lansner, 1998). Jensen et al. investigated a complex

* Corresponding author. Tel.: +49-731-502-4154; fax: +49-731-502-4156.

E-mail addresses: fritz@neuro.informatik.uni-ulm.de (F.T. Sommer), thomas.wennekers@mis.mpg.de (T. Wennekers).

0893-6080/01/$ - see front matter © 2001 Elsevier Science Ltd. All rights reserved.
PII: S0893-6080(01)00064-8

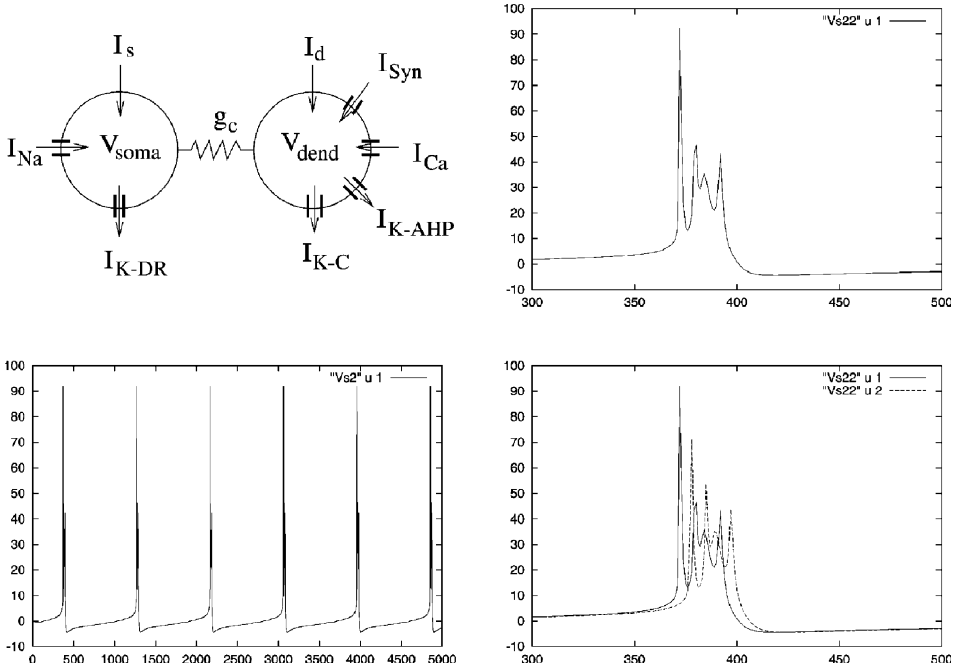

Fig. 1. Qualitative behavior of Pinsky–Rinzel neurons. Upper left figure shows a single neuron scheme consisting of a dendritic and a somatic compartment comprising different active ion-currents and synaptic inputs. A single neuron is capable of firing repetitive complex bursts (lower left and upper right plot). In a network of synaptically connected cells these bursts synchronize, but not the individual spikes constituting the respective bursts: the lower right plot shows soma potentials of two cells out of a total of 100 neurons. (Scales: time is in 0.1 ms, potentials are in millivolts).

multiplexing mechanism of short- and long-term memory using gamma and theta oscillations (Jensen & Lisman, 1996; Jensen, Idiart & Lisman, 1996). Menschik and Finkel (1998) proposed an attractor associative memory as a model for the CA3 region in the hippocampus. In their model, memories are expressed by synchronized gamma activity persisting over a theta cycle.

We combine simplified but biophysically well-tested compartmental-neurons (Pinsky & Rinzel, 1994) with the functional hypothesis of associative long-term memory. Different from the referred previous studies, we constrain the model using quantitative results from studies on more abstract sparse associative networks. The simple and clear-cut functional assumption allows a direct comparison of the biologically motivated model and formal models. In the following simulation experiments, we compare two different situations of external stimulation, persistent and short pulses of current injections in subsets of cells and study the network activation processes. Observing the network activity, ongoing, or induced by stimulation, one is able to address, in particular, the following specific questions.

- In which parameter ranges can the synaptic memory be selectively accessed by stimulation?
- How is spike synchronization and gamma periodicity used for coding?
- Does cell bursting play a role in information coding of memory access?
- How does the computational performance of realistic spiking neurons compare with abstracted neurons?

2. The simulation model

2.1. The model of Pinsky and Rinzel

The computational model used in the following simulation experiments employs the two-compartment neuron model of Pinsky and Rinzel in exactly the same form as described in Pinsky and Rinzel (1994). Because this single neuron model is still quite complex, we will only qualitatively sketch the modeling approach in the following. The exact dynamic equations and parameter settings can be found in Pinsky and Rinzel (1994).

The single neuron model consists of two resistively coupled compartments describing the soma and the dendrites of a nerve cell, respectively (see Fig. 1). The transmembrane voltage V_i of each compartment is given by a current balance equation

$$C\frac{\mathrm{d}V_i}{\mathrm{d}t} = -I_{\mathrm{ionic},i} + g(V_j + V_i). \tag{1}$$

Here, $i, j = 1$ or 2, since we only have two compartments. C is the membrane capacitance and $I_{\mathrm{ionic},i}$ the total ionic membrane current described in more detail below. Furthermore, g is the coupling conductance between soma and dendrite; its value has a strong influence on the spiking behavior of the cell (see Pinsky & Rinzel, 1994, and below).

The total ionic current for each of the compartments consists of a sum of contributions from membrane leakage, different active ion channels, synaptic inputs from other neurons in a network, and inputs from external sources.

Active currents for channel type x in compartment i are in general given by

$$I_{x,i} = g_x m_x n_x (V_i - V_x), \quad (2)$$

where g_x denotes the maximum conductance of a particular type for compartment i and the V_x are reversal potentials for the respective ions. The m_x, n_x are state variables describing activation and inactivation by a first order kinetics (details in Pinsky & Rinzel, 1994).

The soma compartment, similar to the Hodgkin–Huxley model, contains only Na^+ and delayed rectifier K^+ conductances. When isolated, it generates regular action potentials at a frequency determined by the constant external driving currents. The dendritic compartment, on the other hand, contains only Ca^{2+} and calcium mediated K^+ currents (C-type K^+ current, I_{KC}, and after hyperpolarization, I_{KAHP}). Submembrane Ca^{2+} concentration is modeled by a first-order calcium buffer, which controls the maximum conductance of the I_{KC} current. The dendritic compartment is capable of generating slow calcium spikes.

Compared to integrate-and-fire neurons, Pinsky–Rinzel neurons—as compartmental neurons in general—show a much richer dynamic behavior. For instance, the choice of the coupling between the two compartments (cf. the parameter g in Eq. (1)) determines different modes of neuronal activity, single spikes, bursting, or apparently chaotic modes combining spike- and burst-behavior (see Pinsky & Rinzel, 1994, and Fig. 1). In addition, the burst behavior is further controlled by the level of depolarization, and the model undergoes a refractory period after spiking. On a longer time scale of a few hundred milliseconds a firing cell can reveal neural adaptation, well known from electrophysiological recordings in many cortical principal cells.

The model as described in Pinsky and Rinzel (1994) also includes realistic synaptic transmission characteristics, i.e. synaptic time constants, reversal potentials and other properties of AMPA and NMDA synapses. In Pinsky and Rinzel (1994) a biophysical study was performed using a simulated network of 100 cells, randomly connected with a synaptic density of 20% (2000 of the $100(100-1)$ connections possible between cell pairs were randomly chosen and assigned to the same positive value). All cells further received a constant inhibitory soma current. If the soma of a single neuron was stimulated by a brief excitatory current, the network responded with synchronized population bursting in the gamma range that persisted for 400 ms or longer, depending on the maximal NMDA conductance. The NMDA synapses, hence, provided a persistent excitation, but their time constant was too large to provide the phase coupling between the bursting cells. This was due to fast AMPA mediated currents. Blockade of the AMPA synapses led to a desynchronization of the bursts. Thus, the simplified Pinsky–Rinzel network reflects a property that has already been demonstrated in slice experiments in Traub, Wong, Miles and Michelson (1991): the fast AMPA synapses provide a coupling mechanism responsible for burst and spike synchronization.

2.2. Model extensions

We extend the model of Pinsky and Rinzel (1994) to a functional level by storing a set of patterns by Hebbian synaptic modification in a structured connectivity matrix. If not stated otherwise, model parameters of our work comprising $N = 100$ excitatory cells are identical with those in Pinsky and Rinzel's original model. Nonetheless, our model differs with respect to the external stimulation, the structured connectivity matrix, and the kind of inhibition present in the network. In our model excitatory connections are potentially possible between every pair of cells. The actual connectivity depends on the number, P, and size, k, of the stored memory patterns. It is approximately given by $s = 1 - \exp[-Pp^2]$ where $p = k/N$ and ranges from 0.05 to 0.4 in the following experiments. We aim to study a configuration of cortical cells in close vicinity of each other, i.e. within a cortical column. Neuroanatomical studies have estimated a mean density of synaptic contacts of 0.1 between cells with distances lower than the radius of a cortical column (Braitenberg & Schüz, 1991). We also study connectivities higher than 0.1, since local networks will exist with connectivity higher than the mean and since synaptic distribution and memories are likely to be dependent. Memory patterns might not be just determined by afferences and learned by Hebbian synaptic modification, they may well be shaped and selected by the network structure already existing. Cortical subnetworks with high connectivity are the most probable candidates to cooperate in a computational function. We assume that connected cell pairs have synapses in both directions, i.e. that the coupling matrix between the excitatory cells is symmetric. It is not clear to what extent this assumption is idealized. To our knowledge, there are no experimental estimates of the conditional probability of reciprocal connections between a pair of single neurons, given that they have a synaptic contact in one direction. In our model, the synaptic transmission efficacy is formed in a Hebbian learning phase preceding the retrieval trials. The learning is modeled as simplistic as possible. We use the clipped synaptic modification rule of the Willshaw model (Willshaw, Buneman & Longuet-Higgins, 1969), driven by a set of on/off activity configurations (memory patterns) presented to the network. The configurations are random and overlapping, i.e. each contains a fixed number, k, of active cells, and each cell can be active in more than one memory pattern. The specification of the investigated memory tasks (pattern number, P, and sparseness, p) was oriented on the efficient operation range of abstracted sparse associative memories (Palm & Sommer, 1995; Schwenker, Sommer & Palm, 1996; Sommer & Palm, 1999; Willshaw et al., 1969). We used memory patterns with $k = 10$ active neurons. On these tasks, we could directly compare the retrieval performance with theory and simulation results of the original and the feedback Willshaw model.

Since only $k = 10$ cells constitute a memory pattern in our

model, the synaptic strengths, i.e. the AMPA, NMDA and $GABA_A$ conductances, have to be scaled appropriately to yield a meaningful network activity. Other synaptic currents were not modeled. Peak AMPA and NMDA conductances in Pinsky and Rinzel (1994) were $g_{AMPA} = 0.0045$ and $G_{NMDA} = 0.14$ mS/cm^2, respectively. Their work did not consider IPSPs.

Synapses between excitatory neurons terminate on the dendrite-like compartment and activate AMPA- and NMDA-mediated currents as in Pinsky and Rinzel (1994). We run a series of simulations with conductivities scaled up by a factor from 20 to 360 as compared to the values in Pinsky and Rinzel (1994). Thus, the 10 neurons of a memory pattern may be envisaged as representing a cell assembly of 200–3600 cells. The resulting AMPA conductivities are then in the range 0.09–1.62. The NMDA conductances vary proportionally.

The external input into neurons is modeled by Poisson processes with a constant default rate of 500 spikes/s. For technical reasons, the input spikes had an amplitude of 1 and lasted for 0.5 ms. The input processes terminated on the dendritic compartment and exclusively activated AMPA currents. Default conductance was $g^{in}_{AMPA} = 0.9$, but Fig. 7 also employs effects when the input strength is halved or doubled.

In contrast to Pinsky and Rinzel's model, inhibition is not constant in our work. We assume that it is roughly proportional to the instantaneous ensemble averaged firing activity of the principal cells. This assumption is motivated by the fact that in a local network, disynaptic loops from principal cells to interneurons and back should provide an amount of inhibition that depends roughly proportionally on the total network activity. This way, inhibition acts as a threshold control, that moderates the activity level in the network and keeps that activity from unphysiological states where all cells fire at very high rates. This is a traditional but still supported idea about the function of inhibitory interneurons (see, e.g. Nelson & Turrigiano, 1998, for a recent discussion). Further functional roles of interneurons are not excluded. Note that we do not assume that excitation and inhibition are 'balanced', say, in the sense that excitatory and inhibitory currents into pyramidal cells are roughly equal in stationary firing states. In fact, the simulations to be shown in the following reveal that inhibition should be stronger than excitation. This turns out as a condition for high memory capacities, and is in accordance with theoretical works which demonstrate that such a threshold control also improves the memory capacity in abstract associative networks (Hirase & Recce, 1996; Palm & Sommer, 1995; Schwenker, Sommer & Palm, 1999; Wennekers & Palm, 1997, 2000).

Technically, in the simulation model, we do not implement interneurons individually, but instead assume that action potentials of principal cells evoke not only EPSPs on their target cells, but—via inhibitory loops—also IPSPs on all cells in the network. Accordingly, any spike of a principal cell evokes equally weighted IPSCs into all principal cells. These inhibitory synapses employ a fast GABA-ergic conductance change with reversal potential $V_{Cl} = -75$ mV and a shape resembling an alpha function with fast rise-time and slow decay. These conductance changes are actually derived from low-pass filtering of the pyramidal cell spikes through a sequence of three filters with time constants 1, 2 and 7 ms. The maximum conductivity is further varied in the range 1, 2, …, 6 mS/cm^2. (To be compared with our choices of the AMPA conductance: $g_{AMPA} = 0.09$–1.62.)

3. Simulation results

With the simulation model described in Section 2 we performed experiments to study how the stored memories can be recalled (retrieved) by external stimulation. In a retrieval trial, a subset of cells with varying overlap to one of the memory patterns was selected for stimulation. For a defined period of time, the cells in the stimulation subset received depolarizing dendritic input (modeled by Poisson processes).

The stored memories x^1, …, x^P were random configurations of $k = 10$ active cells in the $n = 100$ network. To judge the recall of an addressed memory x^i, we calculated its transinformation between the network state $\hat{x}$ (the binary spike/no spike pattern) and the memory: $t(x^i, \hat{x})$. In the following, we display the *quality* $0 \le Q \le 1$, defined as the retrieved transinformation normalized by the information content of a memory $S = S(x^i) = Ns(k/n)$ with $s(p)$ the Shannon entropy (in our case $S \approx 36$ bits):

$$Q = t\left(x^i, \hat{x}\right)/S \tag{3}$$

For two patterns x, $y \in [0, 1]^N$, the transinformation is defined as $t(x, y) = i(x) - i(x|y)$ where the conditional information is defined as

$$i(x|y) = p_y s(p[x = 0|y = 1]) + (1 - p_y)s(p[x = 1|y = 0])$$

with $p[x|y]$ conditional probabilities and $p_s = |x|/(N - |x|)$ the mean activity of the pattern. The *memory capacity* of the network (in bit/synapse) is given by

$$C = P\, t(x^i, \hat{x})/N^2 = P\, QS/N^2 \tag{4}$$

3.1. Tonic stimulation

In a first series of retrieval experiments, a subset of cells with a defined overlap $l = 5$ to one memory pattern received persistent depolarizing dendritic input. Fig. 2. displays spike trains of two cells during stimulation of five of 10 cells of a stored memory ($N = 100$, $P = 50$, $g_{AMPA} = 0.45$, $g_{Cl} = 3.7$). The upper time course is of one of the directly addressed neurons, i.e. it receives direct afferent input; the second cell displayed receives input from internal connections only. Because it belongs to the addressed memory, the internal

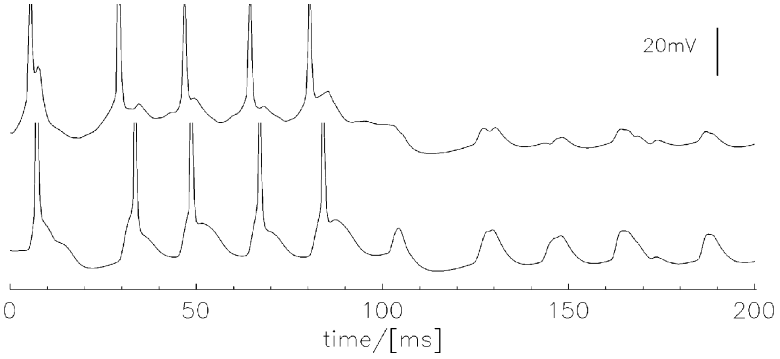

Fig. 2. Typical soma potentials (spikes truncated) for a retrieval trial for $P = 50$, $k = 10$, and five cells of the memory stimulated. The upper cell receives direct external addressing input, the lower cell belongs to the rest of the addressed memory pattern. At time $t = 95$ ms, stimulation switches to another address.

connections provide enough input to bring the membrane potential above firing threshold.

Note that in Fig. 2 the spiking rates are in the gamma range and that the spike activity is synchronized within a time window of a few milliseconds. This is merely a consequence of the synaptic time constants of both the excitatory and inhibitory PSPs. The spikes are phase locked, but the spike of the second cell is delayed due to synaptic integration on the cell membranes. Here, essentially the time constants of excitatory PSPs on the soma matter. On the other hand, the duration of the silent phase after the synchronized population burst is merely determined by the IPSPs evoked by the burst itself. Cell firing is suppressed as long as the inhibition keeps the neurons below firing threshold. So, the frequency of the collective oscillation in this simulation is determined by the relaxation time of inhibitory PSPs, somewhat prolonged by recurrent excitatory transmission steps (cf. also Whittington, Traub & Jeffreys, 1995).

Half time in the sweep displayed in Fig. 2, at time 100 ms, the external stimulation switches to another stimulus pattern. The new stimulus addresses five neurons of a second memory pattern where the two neurons shown in Fig. 2 do not belong to. Notably, the cell responses follow the stimulus withdrawal immediately, that is, there is no afteractivation of the first memory pattern, which would be expressed, for instance, in an ongoing spiking of the neurons (cf. e.g. Jensen et al., 1996). Instead, only subthreshold membrane oscillations remain, which are caused by cross-talk from the new activated memory pattern. Since the memory patterns have mutual overlap, there are potentiated synapses even between neurons not belonging to the same pattern. The membrane potential increase by cross-talk EPSPs is again followed by inhibitory phases. Thus, the rhythmic potential fluctuations in Fig. 2 for times larger than 100 ms consist of superimposed excitatory and inhibitory PSPs. During the first 100 ms, these fluctuations are stronger because the neurons belong to the addressed memory pattern; in fact, there they are strong enough to evoke spiking activity in the respective assembly of memory neurons.

This can be seen more clearly in Fig. 3. The upper frame of Fig. 3 shows a raster plot of the soma potentials for all neurons in the network during a similar retrieval event as in Fig. 2. Neurons 1–10 constitute the memory, the cells 1–5 receive direct stimulation. One can see that directly stimulated cells fire first and the internally driven cells belonging to the same pattern come up a few milliseconds delayed. The temporal synchronization among the internally driven cells is higher since the internal connection convergence and divergence averages over the time dispersion within the set of directly stimulated cells.

The retrieval quality is displayed in the lower frame of Fig. 3. This quality plot quantifies a fact already obvious in the raster plot: at high memory load of $P = 50$ memories,

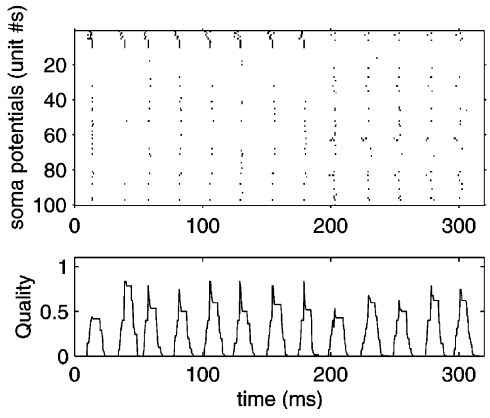

Fig. 3. Retrieval trial with parameters similar to Fig. 2. Stimulation switches at time frame $t = 190$ ms from one memory to another. The upper frame shows a raster plot of soma potentials of all cells. Time runs from left to right and the abscissa counts the neurons. Soma potentials are gray-value coded where black denotes high potentials during spikes. The lower frame displays the retrieval quality trace over time. The quality is determined with respect to the addressed memory consisting of the first 10 cells. Neurons 1–5 receive external stimulation. Note that the network expresses regular spiking in the memory pattern and some spurious action potentials. There are time-locked gamma oscillations, though time-locking is not perfect. The spikes jitter over several milliseconds, mainly caused by excitatory synaptic transmission: address neurons 1–5 always fire first and trigger the cells 6–10. Since for $P = 50$ the memory load is high, retrieval is impaired by cross-talk, i.e. $Q < 1$. The quality is measured continuously for spikes within the previous 10 ms.

the repeated retrieval events yield higher quality than the first completion process just after stimulus onset (at times around $t = 10$ and 110 ms). This is due to the slower neuronal variables that become better aligned after the first population burst. These first bursts after stimulus switch contain considerably more than the 10 neurons of the addressed memory patterns, which lowers the retrieval quality. During later retrieval phases, adaptation processes in the neurons improve the quality, but also during these periods spurious ones appear occasionally. Thus, the pattern recall is cleaned up with respect to spurious activity in the second and further retrieval phases. Furthermore, one observes that the quality often assumes a sharp maximum in the first part of individual retrieval periods. This reflects that the correct cells respond somewhat earlier than spurious cells and could, therefore, be segregated in further processing stages (for instance by coincidence detection).

3.2. Pulsed stimulation

We performed further experiments, where the stimulation was not permanent, but was restricted to a brief period of 12.5 ms. After that, the stimulus was off. To test robustness of the retrieval against background noise, all cells in the network received a permanent background input in form of Poissonian processes at rates 50 ips. Stimulated cells, in addition, received the usual input spike trains of 500 ips.

Fig. 4 shows a retrieval event typical for these series of experiments. Stored are $P = 20$ patterns and five of the neurons of the memory pattern are addressed. This is in the range of medium memory load, where retrieval quality is close to one. The upper and the middle frames show the input spike trains and soma potentials of all cells as raster plots, respectively. The lowest frame reveals the quality of the pattern last addressed by the stimulus. Three stimulation impulses are applied in this run, starting at times 60, 130 and 160 ms, respectively. The cells respond with gamma activity as before. Here, eight cycles are displayed. The first two cycles correspond to ongoing network activity only induced by the background input and recurrent network connections. The third cycle corresponds to a perfect association processes ($Q = 1$). After that, the memory pattern is completely wiped out in the next cycle. Activity in the sixth and seventh cycles corresponds to two subsequent perfect retrieval events triggered by two further stimuli following close to each other. Again, the last cycle shows that the memory pattern is wiped out immediately after the stimulus vanishes. Thus, pulsed stimuli can suppress ongoing network activity quickly, and fast switches between different stimuli are possible at extremely high rates in the gamma range. Note with the example in Fig. 4 that in this retrieval mode the burst occurrence is correlated with retrieval events. Ongoing activity is mostly constituted of regular spikes.

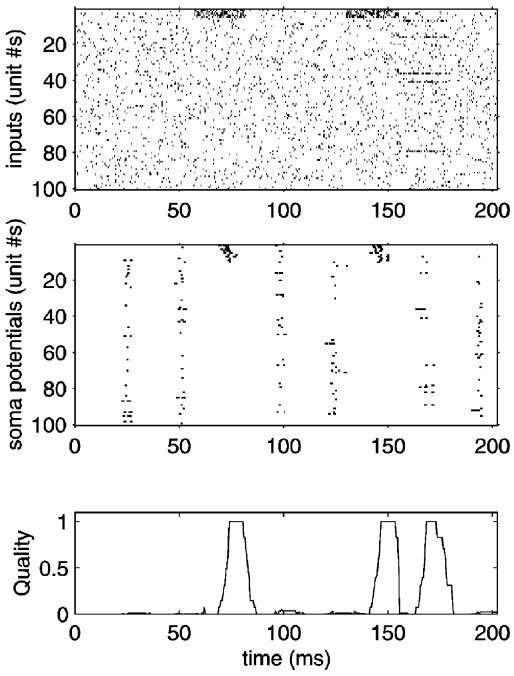

Fig. 4. Raster plot of inputs (upper frame), soma potentials (middle frame) during a retrieval trial with background input and pulsed stimulation. The y-axes display the cell numbers, the x-axes the time, as in Fig. 3. Parameters: $P = 20$; $l = 5$; $N = 100$; $k = 10$; the noise-to-signal ratio is 0.1.

3.3. Variation of model parameters

The simulations in the previous sections demonstrate that the presented biophysically elaborate model is able to retrieve the stored memory patterns very fast and at a high quality, although the memory load is already extremely large. It is known from studies approaching fixed point retrieval in more abstract associative memories that attractor basins become smaller and smaller when the capacity limit is reached, such that eventually a fine-tuned firing threshold is necessary for a proper network operation.

This seems to suggest that also in our network parameters must be fine-tuned to reach high memory loads. In the present section, we demonstrate that the high memory capacity in our experiments is not a matter of critical parameter fine-tuning. The dynamic retrieval model presented here—in contrast to fixed point retrieval—works very well in surprisingly large parameter regimes. This property is mainly based on the temporal structure of the network firing activity, which separates in time the input and correctly retrieved spikes that come early in a gamma period from wrong spurious spikes which on average fire late (cf. Figs. 3 and 4).

Fig. 5 displays results of a large series of experiments where the number of stimulated cells l and the memory load P has been varied systematically. Note that P influences the sparseness of the connectivity matrix (cf. Section 2.2) in the range 0.18 to ≈ 0.5. For high P and low numbers of address bits/neurons the capacity drops because the few

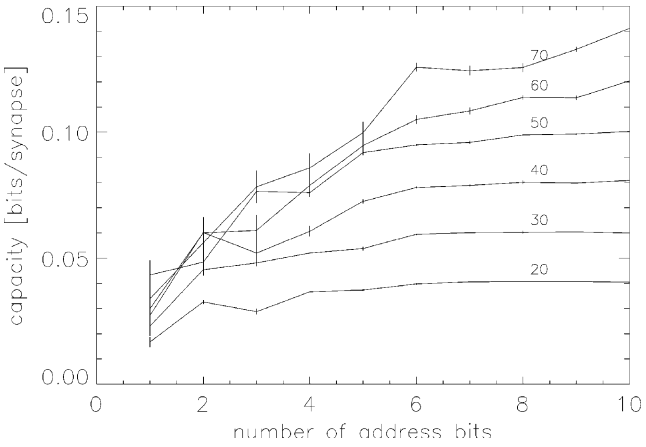

Fig. 5. Memory capacities for different numbers of active address neurons and increasing memory load; $P = 20, 30, \ldots, 70$; $l = 1, 2, \ldots, 10$; $N = 100$; $k = 10$. The capacity is derived from the transinformation in spike patterns, averaged over 10 retrieval periods for 10 randomly selected addresses.

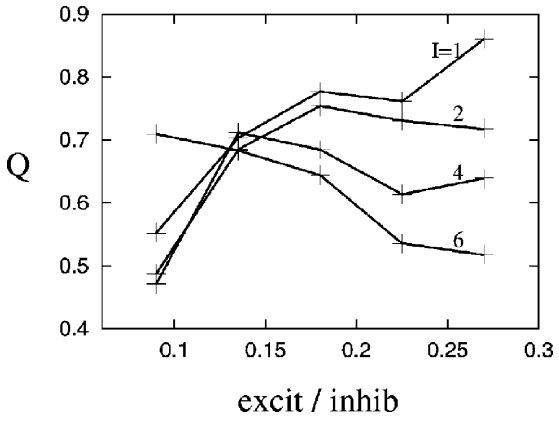

Fig. 6. Average retrieval quality for varied excitatory and inhibitory coupling strengths, and $N = 100$; $k = 10$; $P = 50$; $l = 5$. The abscissa displays the fraction of excitatory to inhibitory peak conductivity. The ordinate denotes the retrieval quality averaged over five retrieval periods for five randomly selected patterns. Different curves correspond with differently strong inhibition.

input bits do not suffice to determine the addressed memory pattern uniquely. Hence, the retrieval quality is impaired.

The longer the plateau in the capacity curve, the better is the input fault tolerance of the recall process. Up to a load of $P = 50$, we find a pronounced input fault tolerance and the corresponding memory capacity is close to 0.1 bit per synapse. This is a striking result since for $P = 50$ the synaptic storage has almost reached the theoretical optimum where 50% of the synapses have been potentiated. The theoretical optimum[1] could be approached even more closely in our model with an optimized memory pattern size k which is lower than $k = 10$ for a network size of $N = 100$ neurons.

To check whether the retrieval performance depends critically on the setting of the EPSP amplitudes and/or the inhibition strength (more precisely, on the excitatory and inhibitory peak conductances), we varied these parameters for a fixed memory load of $P = 50$ and a constant number of stimulated cells of $l = 5$. Other parameters were as stated earlier. Fig. 6 displays quality values averaged over five retrieval periods for five randomly selected memory patterns. Different curves represent different levels of inhibition, $g_{GABA} = 1, 2, 4, 6$, and the abscissa denotes the relative level of AMPA- to GABA-conductance (NMDA) is proportional to AMPA, see Section 2.2). The quality measure is directly proportional to memory capacities by means of Eq. (4).

Note that the abscissa in Fig. 6 is a relative scale. Thus, if we increase the inhibition, we have to increase excitation simultaneously to reach a good retrieval quality. This, however, does not mean that excitation and inhibition are 'balanced'. In fact, in the simulations in Fig. 6 the inhibitory conductances were always at least three times larger than the excitatory ones.

[1] For finite size networks, the theoretical optimum is below the asymptotic limit of 0.69 bit per synapse. For instance, with a size of 512 cells the optimum is 0.4 bit per synapse (see Palm & Sommer, 1995).

Moreover, the figure displays much of the whole range of conductances where we were able to reach a significant retrieval quality at all. This range is restricted essentially by two constraints which become relevant at low and high values for the relative coupling strengths, respectively. For instance, if the ratio between excitation and inhibition is much smaller than one, i.e. if the inhibition is very much stronger than excitation, then the network progressively loses the capability to complete the input patterns at all. This is because the few addressed neurons already charge the inhibition so much that the non-addressed cells can hardly reach threshold, driven only by the weak recurrent signals. Those fire only in some gamma periods, if the excitation is strong enough. Then, the quality can be rather good, as in the plot for $I = 1$, and the excitation equal to 0.9, but retrieval events are rare (roughly each fourth gamma period for the parameters shown). On the other hand, if the excitatory coupling strength is low, the addressed cells fire almost asynchronously. This also leads to only a few retrieval events, but in addition to a considerable amount of incomplete retrievals and wrong ones. Therefore, the quality in Fig. 6 is bad if both excitation and inhibition are small.

If the coupling strengths increase, the capacity obviously displays quite a large region where it is relatively high. However, for weak inhibition the curves increase, whereas they fall for strong inhibition. The increase for weak inhibition is a consequence of the fact that larger synaptic strengths support more pronounced gamma oscillations. So, for weak inhibition, e.g. $I = 1$, and weak excitation the activity is almost asynchronous and retrieval is only seldom and bad; but if excitation increases, the network starts to oscillate, which leads to a better aligned timing of the cells, just as displayed in Fig. 3. First, the externally addressed cells fire, then the missing ones of the addressed pattern, and finally some spurious ones. As long as the coupling strengths are not

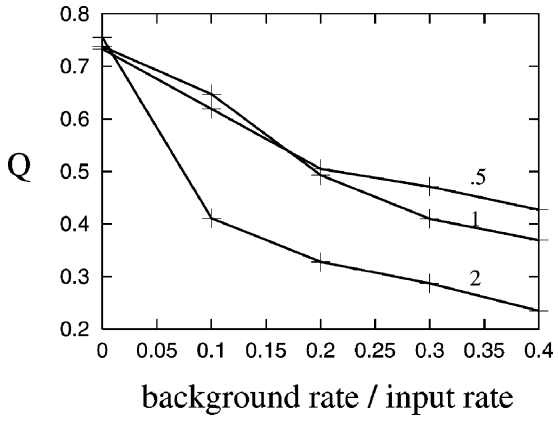

Fig. 7. Average retrieval quality if the input strength and background firing are varied; $N = 100$; $k = 10$; $P = 50$; $l = 5$. The input firing rate is fixed at 500 ips. The abscissa denotes the fraction of the background rate and the input rate. The quality is averaged over 10 retrieval periods for 10 randomly selected patterns. Different curves are for different input strengths (default: 1; halved: 0.5; doubled: 2).

too strong, this dynamic retrieval mechanism increases the memory capacity as long as excitation grows.

Nonetheless, if the coupling strengths increase further, that is, if the excitation-to-inhibition ratio is larger than 0.25 then a different source for retrieval errors becomes dominant. In this case, the dynamics gets faster and faster such that even suboptimally driven neurons reach their firing threshold so quickly that the somewhat slower inhibition is no longer able to suppress or at least delay the firing of these wrong neurons. Therefore, the oscillation can no longer separate the correct and spurious spikes in time. Thus, the retrieval events are "poisoned" by many spurious ones, which reduce the average retrieval quality. This effect is most clearly seen in the decreasing curve for $I = 6$ in Fig. 6.

All in all, Fig. 6 demonstrates that good retrieval qualities can be obtained in large parameter domains for the coupling strengths. Fig. 7 shows, that the network is also quite robust against input noise. Whereas Figs. 5 and 6 are obtained without background noise, in Fig. 7 retrieval qualities are displayed when all neurons receive a certain amount of background activity, as can be seen in Fig. 4. This background noise is modeled as Poissonian input with a fixed firing rate into all cells. The abscissa in Fig. 7 measures the 'noise-to-signal ratio', i.e. the fraction of the background rate (500 ips) and the input firing rate. Individual curves in the figure denote different levels of absolute input strength. The default value, as stated in Section 2.2, is $g^{\text{in}}_{\text{AMPA}} = 0.9$. The three curves are for the default value, as well as halved and doubled input strengths. Changes in input strength influence both the background as well as the input fibers. The number of stored patterns is $P = 50$ and the number of addressed neurons $l = 5$.

As can be read from the figure, the quality is insensitive to changes in input strengths in the simulated range as long as there is no background activity, $Q = 0.75$ for zero background firing. As one would expect, background noise deteriorates the retrieval quality. However, the decrease in quality is only about 10–15% for a signal-to-noise ratio lower than 0.1 and an input strength of 0.5–1.5. Thus, there is some robustness against input noise at high memory load. In the high quality retrieval regime at medium load ($P = 20$), a noise-to-signal ratio of 0.1 hardly impairs the retrieval results (for an example, see again Fig. 4).

Strikingly, one observes in Fig. 7 that stronger (doubled) input strength lowers the memory capacity. The reason for this is that the firing rate of the background neurons depends supralinearly on the input noise (threshold effect, i.e. sigmoid rate function). Strong background noise, therefore, leads to considerably higher spontaneous firing rates of unaddressed neurons, against which the firing rates of the externally addressed neurons become progressively negligible.

Similarly, though not shown in the figure, a very weak input strength also leads to only small quality values. This is because in that case the network activity decouples from the input and becomes more and more determined by network intrinsic processes such as spontaneous firings of cells and the recurrent distribution of spikes to other cells.

4. Results and conclusions

This paper examines the computational role of the spatial and temporal fine structure of neuronal activity, in particular, spiking, bursting, synchronization and oscillations in circuits of strongly connected cortical cells. Our modeling approach assumes associative memory function and uses model constraints from studies on abstract associative memory.

4.1. The modeling approach

We propose a computational model for a group of cortical cells employing the quite faithful biophysical description of neurons and synapses by Pinsky and Rinzel and a rather simplified version of Hebbian associative synaptic learning. The model design and parameter settings are based on a combination of results obtained by different biophysical and associative network studies. The latter examine a computational function likely in cortical networks at a high degree of biophysical abstraction. They introduce and assess the criterion of performance efficiency of a computational model. In our model, we assume that the hypothetical computational function uses the underlying biophysical mechanisms efficiently. This pins down a number of additional modeling features: sparse memory patterns and, since the sparse memory performance does not strongly depend on the details of synaptic learning, the strong simplification of the learning process. The efficiency assumption lacks direct empirical justification, but one can argue that it is implicitly inherent in the functional hypothesis providing its universality even under perturbations neglected in the model. Furthermore, evolution theory

argues for an optimization of biological functions during ontogenesis. No a priori assumptions are required in our model about the nature of temporal coding. In the simulation experiments, we just observe the network activity during and after external stimulation and compare it with the spatial patterns that have been previously stored by Hebbian synaptic potentiation.

4.2. Interpretation of the simulation results

In our model, the network activity travels in avalanches through associative excitatory connections from highly excited (say afferently driven) cells to less excited ones. This excitation is mediated through the structured excitatory connections. In cells with strong postsynaptic input, the AMPA-mediates PSPs cause the spiking via selective associative connections. This way, starting from afferently driven cells, activity can propagate in several mono-, di-,and tri-synaptic associative steps. Due to synaptic transmission times, one, therefore, finds synchronization only within a few milliseconds between spike events in different cells. In addition, the larger the set of spiking cells in the network, the stronger inhibition builds up until it finally suppresses further spiking. Nonetheless, if afferent excitation persists, or also due to the slow NMDA synaptic transmission over excitatory connections, after decay of the inhibition, another population burst appears. We observe this rhythmic network activity to lie in the gamma frequency range (20–80 Hz).

We examined how memories stored in the connectivity structure of the network can be recalled by external stimulation. Among the variety of possible spatio-temporal stimulation patterns, we chose two for closer inspection by simulation experiments.

First, we studied memory retrieval provided by tonic stimulation. As had been shown in similar network simulation studies, for instance in Menschik and Finkel (1998), a pattern of synchronized spike events represents properties of the addressed memory. All cells in the externally addressed memory, whether or not directly stimulated, participate in phase-locked gamma activity. At high memory load, the association in the first gamma cycle after stimulus onset already provides a reasonable estimation of a memory pattern, but contains some erroneous responses of non-addressed cells. The second synchronized spike event, i.e. the second gamma cycle after stimulus onset, provides the memory with increased quality. The improvement is due to the alignment of the slow neuronal variables (neuronal adaptation) that requires time of more than one gamma cycle. The quality in the cycles after the second stays almost constant and the time lag between primary and secondary driven cells stays constant in all cycles. If the stimulus is changed to another spatial pattern, the network switches within one gamma cycle to the new pattern.

Second, we examined memory retrieval evoked by a short stimulus pulse (12.5 ms). The main difference to persistent stimulation is that memories are not attractors of the network dynamics. The network responds to the stimulus with only a single association event occurring in the gamma period directly following the stimulus onset. In the load domain providing almost error free retrieval with persistent stimulation, pulsed stimulation yields almost the same performance. As can be observed in Fig. 4, pulsed retrieval is also robust with respect to ongoing network activity and background input activity. The retrieval is still error free at a load $P = 20$, a memory load where the capacity curve for noise free retrieval with persistent stimulus in Fig. 4 already slightly decreases due to a (small) finite retrieval error rate. Fig. 4 also demonstrates that subsequent stimulation pulses addressing different memories can achieve readout rates up to the gamma frequency. The maximum readout speed of the network is limited by the gamma rhythm of the network. We observed that with pulsed stimulation the expression of bursts and regular spikes became correlated with the retrieval process. In Fig. 4 most bursts occur during pattern retrieval.

The robustness of the demonstrated associative memory function in the simulation network was explored by retrieval quality measurements under the variation of different parameters. High information capacity is expressed in a wide range of excitation/inhibition ratios (Fig. 6). At high memory load the quality decreases, with graceful degradation, not abruptly (Fig. 7), at medium load there is almost no impairment up to noise-to-signal ratios of 0.1 (Fig. 4).

The presented experiments explore a parameter range where the internal wiring provides basically a feedforward completion of partially addressed memory patterns. Without persistent bias, the memories are not attractors. Thus, recurrent activity flow that can improve the retrieval quality iteratively, as suggested in Amit (1995) for local networks in area IT, plays no major role in our experiments. However, they appear in a pair of Pinsky–Rinzel neuron pools that are bidirectionally coupled with much weaker connectivity strength (Sommer & Wennekers, 2000, 2001). There, recurrent reverberations enhance the associative memory operation.

4.3. Conclusions

As a general result, our experiments demonstrate that biologically realistic networks provide in a wide parameter range robust associative storage of sparse patterns at a capacity close to the one of technical networks. In the light of our model, gamma oscillations can indicate sequences of fast individual retrieval processes. However, gamma oscillations are also found in the background activity, and with flashed stimulation retrieved memories are not coded by a rhythmic activity. This suggests to consider the elicitation of a single synchronized population burst (within one gamma cycle) as an elementary association or retrieval process. Gamma oscillations, then, appear as sequences of fast individual retrieval processes carried by associative excitatory connections, rhythmically interrupted by inhibitory interneurons, as has already been proposed based on simulation

experiments in simpler models. This interpretation of gamma activity as has already been proposed earlier (Wennekers, Sommer & Palm, 1995). It stands in contrast to the assumption of coding by periodically firing neurons that simply adjust their phases in time such that different objects become segregated into different phases of the collective oscillations (cf. the discussion in Wennekers & Palm, 1997). Our findings favor functional interpretations that avoid the strong assumptions of phase-coding and strictly rhythmic firing of single neurons.

Our experiments with flashed stimuli revealed that under such stimulation conditions bursts become relevant for information coding. By fast switching between flashed stimuli, a sequence of retrieval processes can be evoked. The maximum readout speed is only limited by the gamma frequency of the network. At maximum readout speed, every gamma cycle can display a different memory.

References

Amit, D. J. (1995). The Hebbian paradigm reintegrated: local reverberations as internal representations. *Behavioural and Brain Sciences*, *18*, 617–657.

Braitenberg, V., & Schüz, A. (1991). *The anatomy of the cortex. Statistics and geometry*, Berlin, Heidelberg, New York: Springer.

Fransén, E., & Lansner, A. (1998). A model of cortical associative memory based on a horizontal network of connected columns. *Network*, *9* (2), 235–264.

Hirase, H., & Recce, M. (1996). A search for the optimal thresholding sequence in an associative memory. *Network*, *7* (4), 741–756.

Jensen, O., & Lisman, J. E. (1996). Theta/gamma networks with slow NMDA channels learn sequences and encode episodic memory: role of NMDA channels in recall. *Learning & Memory*, *3*, 264–278.

Jensen, O., Idiart, A. P., & Lisman, J. E. (1996). Physiologically realistic formation of autoassociative memory in networks with theta/gamma oscillations: role of fast NMDA channels. *Learning & Memory*, *3*, 243–256.

Lansner, A., & Fransén, E. (1992). Modelling Hebbian cell assemblies comprised of cortical neurons. *Network*, *3*, 105–119.

Menschik, E. D., & Finkel, L. H. (1998). Neuromodulatory control of hippocampal function: towards a model of Alzheimer's disease. *Artificial Intelligence in Medicine*, *13*, 99–121.

Nelson, S. B., & Turrigiano, G. G. (1998). Synaptic depression: a key player in the cortical balancing act. *Nature Neuroscience*, *1* (7), 539–541.

Palm, G., & Sommer, F. T. (1995). Associative data storage and retrieval in neural networks. In E. Domany & J. L. van Hemmen, *Models of neural networks III* (pp. 79–118). New York: Springer.

Pinsky, P. F., & Rinzel, J. (1994). Intrinsic and network rhythmogenesis in a reduced Traub model for CA3 neurons. *Journal of Computational Neuroscience*, *1*, 39–60.

Schwenker, F., Sommer, F. T., & Palm, G. (1996). Iterative retrieval of sparsely coded associative memory patterns. *Neural Networks*, *9* (3), 445–455.

Sommer, F. T., & Palm, G. (1999). Improved bi-directional retrieval of sparse patterns stored by Hebbian learning. *Neural Networks*, *12*, 281–297.

Sommer, F. T., & Wennekers, T. (2000). Modeling studies on the computational function of fast temporal structure in cortical circuit activity. *Journal of Physiology—Paris*, *94* (5/6), 473–488.

Sommer, F. T., & Wennekers, T. (2001). Associative memory in a pair of cortical cell groups with reciprocal connections, *Neurocomputing* (to appear).

Traub, R. D., Wong, R. K. S., Miles, R. & Michelson, H. A. (1991). A model of a CA3 hippocampal pyramidal neuron in cooperating voltage-clamp data in intrinsic conductances. *Journal of Neurophysiology*, *66*, 635–650.

Wennekers, T., & Palm, G. (1997). On the relation between neural modeling and experimental neuroscience. *Theory in Biosciences*, *116*, 273–289.

Wennekers, T., & Palm, G. (2000). Cell assemblies, associative memory and temporal structure in brain signals. In R. Miller (Ed.), *Time and the brain. Conceptual advances in brain research* (vol. 3). Harwood Academic Publishers, Amsterdam, pp. 257-273.

Wennekers, T., Sommer, F. T., & Palm, G. (1995). Iterative retrieval in associative memories by threshold control of different neural models. *Supercomputers in brain research* (pp. 301–319). .

Whittington, M. A., Traub, R. D., & Jeffreys, J. G. R. (1995). Synchronized oscillations in interneuron networks driven by metabotropic glutamate receptor activation. *Nature*, *373*, 612–615.

Willshaw, D. J., Buneman, O. P., & Longuet-Higgins, H. C. (1969). Non-holographic associative memory. *Nature*, *222*, 960–962.