

Why spikes? Hebbian learning and retrieval of time-resolved excitation patterns

Wulfram Gerstner, Raphael Ritz, J. Leo van Hemmen

Physik-Department der TU München, James-Frank-Strasse, D-85747 Garching bei München, Germany

Received: 24 October 1992/Accepted in revised form: 21 April 1993

Abstract. Hebbian learning allows a network of spiking neurons to store and retrieve spatio-temporal patterns with a time resolution of 1 ms, despite the long postsynaptic and dendritic integration times. To show this, we introduce and analyze a model of spiking neurons, the spike response model, with a realistic distribution of axonal delays and with realistic postsynaptic potentials. Learning is performed by a local Hebbian rule which is based on the synchronism of presynaptic neurotransmitter release and some short-acting postsynaptic process. The time window of this synchronism determines the temporal resolution of pattern retrieval, which can be initiated by applying a short external stimulus pattern. Furthermore, a rate quantization is found in dependence upon the threshold value of the neurons, i.e., in a given time a pattern runs n times as often as learned, where n is a positive integer ($n \geq 0$). We show that all information about the spike pattern is lost if only mean firing rates (temporal average) or ensemble activities (spatial average) are considered. An average over several retrieval runs in order to generate a post-stimulus time histogram may also deteriorate the signal. The full information on a pattern is contained in the spike raster of a single run. Our results stress the importance, and advantage, of coding by spatio-temporal spike patterns instead of firing rates and average ensemble activity. The implications regarding modelling and experimental data analysis are discussed.

1 Introduction

One of the most important and intriguing problems of neural network theory and neural data interpretation concerns the coding procedure which is used by a biological system to map data from the outside world onto internal network states. Most model approaches to neural networks as well as many experiments on biological networks have been based on the idea that it is the mean firing rate which carries the essential details. This is the tacit assumption that has to be made if a single neuron is

modelled through an 'analog' or 'graded-response' unit, as in most backpropagation networks (Rumelhart and McClelland 1986) and in many other model networks (Hopfield 1984; Kühn et al. 1991). The same assumption is often made when the neuron is described by a two-state spin variable as in the standard model of associative memory (Hopfield 1982; Amit et al. 1985, 1987). In both cases, the output of the model neuron can be interpreted as the mean firing rate (number of spikes per time interval) of an ideal biological neuron.

The above model parallels traditional experiments where neurons have been studied and classified based mainly on their mean firing rate (Adrian 1926; Hodgkin 1948). The experimental advantage of such an approach lies in the fact that mean firing rates are comparatively simple to measure. Once an electrode is positioned next to a neuron (which may require considerable experimental effort), one simply counts the number of action potentials and divides by the measurement time T . Depending on the external stimulus the firing rate can be high or low. If the experiment is repeated while different neurons are recorded, an excitation pattern in terms of high and low firing rates can be measured. This pattern is often interpreted as the answer of the biological network to the specific stimulus applied in the experiment. See Miyashita (1988) for one recent experiment along these lines.

Such an approach, however, may be too simplistic. Since the determination of the mean firing rate requires an average over a sufficiently long time interval (typically 100 ms and more), all information about the exact firing times of a single neuron is lost. The firing time of a neuron compared with the start of a stimulus or the spikes of other neurons could, however, be of great importance for correct interpretation of the stimulus (Perkel et al. 1967; Gerstein and Perkel 1972; Abeles 1982; Aertsen et al. 1986; Johannesma et al. 1986; Palm et al. 1988). Indeed, it has been shown in experiments on the fly that the timing of a single spike becomes important if one attempts to decode the spike train and reconstruct the stimulus (Bialek et al. 1991). To phrase it differently, a fly would be completely lost in a natural environment if the mean firing rate were the only available carrier of information. Similar results can be achieved by a theoretic analysis of the maximum information content

of neuronal spike trains. A coding scheme based on single spikes or interspike intervals yields a much higher information capacity than a rate code (MacKay and McCulloch 1952; Stein 1967; Abeles and Lass 1975; Eckhorn et al. 1976).

To achieve a better time resolution while keeping a good signal-to-noise ratio, two different approaches have been adopted. The first is based on the assumption that there are large groups of neurons, e.g., columns in the visual cortex, which are involved in the same signal-processing task (Hubel and Wiesel 1962). If this were so, it would be enough to measure the mean or mass action of such a group of neurons (see, e.g., Abeles 1982). This can be done with a single extracellular electrode tuned to the global 'background' potential of neuronal activity. If the group is sufficiently large and if all neurons are doing the same job, then the ensemble activity can also be taken as a measure of the typical answer of a single neuron in terms of spiking probabilities. If, however, neurons in this group have different intrinsic properties or different afferent connections or a nontrivial interconnectedness, then all information about individual neurons is lost by taking an ensemble average.

The second time-resolved method is the so-called post-stimulus time histogram (PSTH) (Perkel et al. 1967). After the start of the stimulus, the spikes of one or a small set of neurons are recorded. In a PSTH, the response is accumulated over several repetitions of an experiment with the same stimulus pattern. Such an approach, however, is useful only if the experimental setup, in particular the internal state of the network, is exactly the same in each run. Since in most biological systems the internal state of the network cannot be controlled, this condition is hard to fulfill under realistic circumstances. If we cannot assure an exact repetition of the experiment, averaging over several runs may spoil information which is contained in any single run. That is why we have to take the results of a single run seriously.

To sum up the above consideration, we have argued that the three traditional approaches – averaging over time to get the mean firing rate, averaging over space to get the ensemble average, and averaging over several runs to accumulate a PSTH – may miss valuable information which is contained in the exact spiking times of neurons in a large network. Recent developments in multi-electrode recordings (Gerstein et al. 1983; Krüger 1983; Reitboeck 1983) make it possible to circumvent the above disadvantage and to measure the spikes of several neurons in a single run (see, e.g., Eckhorn et al. 1988; Krüger and Aiple 1988). The results are stored in terms of a spike raster which can be used for careful evaluation thereafter.

In the present study we would like to catch up with these developments from a theoretical point of view. In our model of a neural network, the *spikes of single neurons* are in the center of interest. Each neuron is modelled by a threshold-fire unit with a spike-after-potential equivalent to an increase of a dynamic threshold. Transmission of the spikes along the axon to the synapse is taken into account by a delay Δ^{ax} that depends on the distance between the soma and the synapse. If

a spike arrives at the synapse, an excitatory or inhibitory postsynaptic potential (EPSP or IPSP) is induced, which is also described in a realistic and time-resolved manner. A single EPSP or IPSP can be understood as a neuron's response to an incoming spike. Similarly, the spike-after-potential or dynamic threshold is the neuron's internal response to spike emission. Since both effects are included, the model is called the *spike response model*.

The main idea of our approach is that the answer of a neural system to an external stimulus consists in a complex spatio-temporal spike pattern which is not necessarily associated with a change of mean firing rates or ensemble activities. Such a spike pattern could, in principle, change on a time scale of a few milliseconds, which is quite fast if it is compared with the mean firing rate of real neurons. Under these circumstances, it would be necessary to take a single spike seriously. If we neglect a temporal resolution below a single spike width, we can take an interval of 1 ms as the basic time step of neuronal signal processing in our model.

In view of the much longer synaptic and dendritic integration time constants (typically 5–50 ms) the following questions appear. Is it really necessary to construct a model with a time resolution of 1 ms? Is it at all possible to store time-resolved excitation patterns in the synaptic efficacies despite the long postsynaptic integration time? Furthermore, is it possible to "learn" spatio-temporal activation patterns by a "Hebbian" learning rule? The answer to these questions is affirmative and will be given in Sect. 2.

If this is done, we can address the central question outlined above: How is the information on the stimulus and, more generally, on the outside world encoded in the network? Can it be found in the mean firing rates, ensemble activities, or in the PSTH, or do we need the exact firing time of each spike of every neuron? Using some generic examples we show in Sect. 3 that complex spatio-temporal spike patterns require the full information contained in the spike raster of a single run. If only mean firing rates, ensemble activities or PSTHs are considered, information contained in the correlations of the spikes of different neurons can be missed. If spike rasters of several runs of the experiment are available, they can be used to determine the averaged cross-correlation between pairs of neurons explicitly. In any case, averaging – be it over space, time, or repeated runs – should come at the end of the evaluation procedure and should not be a tacit assumption underlying models or experiments. Conclusions from the results of Sect. 3 are drawn in Sect. 4.

2 Model network

Our network (Fig. 1) consists of N model neurons which are fully connected via asymmetric "Hebbian" synapses. The neurons communicate with each other through the exchange of spikes which are transmitted along the axons. If a spike passes a synapse, an EPSP is evoked at the soma of the receiving neuron. Furthermore, spike




Fig. 1. Network structure. Each pair of neurons in the network layer is connected by several synapses located at different positions along the axon of the presynaptic neuron. In addition to the internal synaptic contributions, each neuron receives an external stimulus from the receptor layer and a common signal from the background of the network. One inhibitory neuron controls the activity of the whole network




Fig. 2. Excitatory postsynaptic potential (EPSP). The typical response $\varepsilon(s)$ at the soma of a postsynaptic neuron is plotted as a function of time after firing of a presynaptic spike. There is no signal for the first few milliseconds due to the axonal delay $1 \leq \Delta^{ax} \leq 4$ ms (here $\Delta^{ax} = 2$ ms). The rise time of the response is $\tau_r = 3$ ms, the overall duration (width of half-maximum) approximately 8 ms. The response function $\varepsilon(s)$ is normalized to 1 and has to be weighted with the synaptic efficacy factor J_{ij}^{ax} to get the EPSP

emission results in an increased firing threshold at the soma of the sender neuron (spike response model, SRM).

Spikes of neuron j are considered as point-like events and described by a Dirac δ -function $\delta(t - t_j^f)$ where t_j^f is the firing time of the f^{th} spike of neuron j . The EPSP evoked at neuron i is described by a realistic response function which extends over several milliseconds (Fig. 2). In our discrete-time model, the EPSP is quantized in steps of 1 ms. The EPSPs of all signals arriving at neuron i are added linearly and yield the total postsynaptic potential $h_i^{\text{syn}}(t)$. For the sake of convenience only, a single inhibitory neuron which is connected to all neurons controls the total activity of the network through a continuous inhibitory signal $h_i^{\text{inh}}(t)$. In reality inhibition is local. In addition to these internal contributions, each neuron receives action potentials from a receptor layer through an input channel $h_i^{\text{ext}}(t)$. All neurons also receive a common slowly varying background signal $h_i^{\text{back}}(t)$ which is the same for all neurons. It describes an average signal from other parts of the network which are not modelled explicitly, e.g., from higher areas involved in association and complex signal processing tasks. The

various elements of our model system are analyzed in more detail in the following subsections.

2.1 Model of spike generation in a single neuron

The description of a single model neuron i is based on three classical notions of neurobiology, namely, firing threshold, refractoriness, and noise. The model has been discussed in extenso by Gerstner and van Hemmen (1992a) and will be reviewed briefly here. In a noiseless idealization, a spike would be induced as soon as the membrane potential $h_i(t)$ exceeds a threshold θ . But due to omnipresent noise, the actual spike may occur a little too early or too late compared with the formal threshold crossing time. To take this into account, we introduce the probability P_F of firing during a short time interval between t and $t + \Delta t$ given a membrane potential $h_i(t)$:

$$P_F(\Delta t | h_i) = 1 - \exp[-\Delta t / \tau(h_i)] \approx \Delta t / \tau(h_i) \quad \text{for } \Delta t \ll \tau(h_i) \quad (1)$$

The response constant $\tau(h_i) = \lim_{\Delta t \rightarrow 0} [\Delta t / P_F(\Delta t | h_i)]$ gives the mean delay of a spike – averaged over repeated runs. As a model assumption we take in analogy to the Arrhenius law

$$\tau(h_i) = \tau_0 \exp[-\beta(h_i - \theta)] \quad (2)$$

where τ_0 is the response constant at threshold. In the limit of a noiseless neuron ($\beta \rightarrow \infty$), (2) reduces to

$$\tau(h_i) = \begin{cases} \infty & \text{for } h_i < \theta \\ \tau_0 & \text{for } h_i = \theta \\ 0 & \text{for } h_i > \theta \end{cases} \quad (3)$$

Thus neuron i stays quiescent if $h_i < \theta$ ($\tau = \infty$) and fires immediately if $h_i > \theta$ ($\tau = 0$).

Whenever a real neuron fires, a complicated sequence of channel opening and closing is induced at the membrane of its axon as well as of its dendritic tree. The net result is a sharp spike and various spike after-potentials. Some of these potentials, e.g., the late after-hyperpolarization (AHP), have very long time constants of up to 100 ms and more (Lancaster and Adams 1986). In order to describe possible spike after-potentials or, accordingly, dynamic increases in threshold in a qualitative model we add after each spike a refractory potential

$$h_i^{\text{ref}}(t) = \sum_{f=1}^F \int_0^{\infty} d\tau \eta(\tau) \delta(t - \tau - t_i^f) \quad (4)$$

where t_i^f with $f = 1, \dots, F$ are the most recent firing times of neuron i . F is a useful parameter that allows switching between adaptive and non-adaptive behavior. For $F = 1$ firing is a mathematical renewal process, since only the last spike contributes. For $F > 1$ the model exhibits adaptive behavior.

The refractory function $\eta(\tau)$ describes a neuron's response to spike emission and is given by

$$\eta(\tau) = \begin{cases} -\infty & \text{for } 0 < \tau \leq \tau_{ref} \\ -\eta^{AHP}/(\tau - \tau_{ref}) & \text{for } \tau_{ref} < \tau < \tau_{max} \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

It has two effects. First, the infinite negative contribution for times $\tau < \tau_{ref}$ (where τ_{ref} is the absolute refractory period) is a formal way to model refractoriness. The strict inhibition prevents the immediate emission of another spike during a time τ_{ref} if a spike has been fired at $\tau = 0$. Second, for $\tau > \tau_{ref}$ we include a slowly decaying inhibitory tail. This is a convenient way to model post-spike refractoriness as being due to K inactivation. The main effect of refractoriness or AHP is a higher effective threshold for the firing of a second and a third spike following a first spike.

In the simulations presented in Sect. 3, we take an absolute refractory time of $\tau_{ref} = 3$ ms, an AHP potential of $\eta^{AHP} = 3$ input units (i.u.) and a threshold $\theta = 0.20$ i.u. The AHP ends at $\tau_{max} = 100$ ms. Unless stated otherwise the noise parameter is $\beta = 12$. The gain function of a model neuron with these parameters exhibits a "soft" threshold at $h \approx 0.5$ i.u. and a maximum firing rate of 330 Hz. A thorough discussion of the model neuron that includes typical spike trains and an exact expression for the mean firing rate at various input levels can be found in Gerstner and van Hemmen (1992a) and Gerstner (1993). For the following it is sufficient to note that such a model neuron has a reasonable gain function and is able to generate realistic spike trains.

2.2 Synaptic transmission

In our network each pair of neurons is connected by several synapses (Creutzfeldt 1983). They are located at different positions along the axon of the presynaptic neuron. If one of the neurons – say neuron j – fires, a spike is transmitted along the axon and induces – after a delay Δ^{ax} – a postsynaptic signal at the receiving neuron i . The delay Δ^{ax} depends on the position of the specific synapse on the axon and varies in our simulation between $\Delta_{min} = 1$ ms and $\Delta_{max} = 4$ ms. To phrase it differently: Given that neuron j fires, neuron i is affected after 1, 2, 3, and 4 ms but the synaptic efficacies may differ for the four delays. This may seem a rather implausible assumption since synaptic contacts of a single neuron to another neuron are unlikely to span time differences as large as 4 ms. The model can, however, be changed to the diluted case where three of the four connections are cut at random. The remaining connection has only one delay with $1 \leq \Delta^{ax} \leq 4$ ms, which seems to be a realistic range. The important point is that the synaptic efficacy of a connection from j to i depends not only on j and i but also on the delay; this will be explained below.

Whereas the axonal spike is a soliton-like event, sharply located in space and time, the response of the postsynaptic neuron is smeared out and gives a much broader signal at the soma, the EPSP. The EPSP has been repeatedly measured in experiments and can be

described by an α -function (Jack et al. 1975)

$$\varepsilon(\tau) = \frac{\tau}{\tau_s} \exp\left(1 - \frac{\tau}{\tau_s}\right) \quad (6)$$

which is shown in Fig. 2. It is normalized to a maximum value of 1 i.u. For the time constant we take $\tau_s = 3$ ms. Thus, in our discrete-time model the rising time encompasses three time steps.

Whereas the time course $\varepsilon(\tau)$ of the EPSP is taken to be the same for all neurons and synapses, the amplitude of the response is specific for a single synapse. It depends on the pair (i, j) of neurons as well as on the position of the synapse on the axon as indicated by the delay Δ^{ax} and is described by a synaptic efficacy factor J_{ij}^{ax} . How the synaptic efficacy can be adjusted by a Hebbian learning rule will be discussed in the next subsection. At the moment we simply assume that all J_{ij}^{ax} are known and kept constant.

The synaptic signals received by neuron i can be added so as to find the total postsynaptic potential

$$h_i^{syn}(t) = \sum_{j=1}^N \sum_{\Delta^{ax}=\Delta_{min}}^{\Delta_{max}} \sum_{f=1}^{n_f} J_{ij}^{ax} \int_0^{\infty} ds \varepsilon(s) \delta(t - s - \Delta^{ax} - t_f^j) \quad (7)$$

For the sake of simplicity, we have assumed in (7) that all contributions add up linearly. It is straightforward to allow for more complicated interactions. In analogy to the excitatory synaptic contribution we model the inhibitory signal by

$$h_i^{inh}(t) = - \sum_{j=1}^N \sum_{\Delta^{ax}=\Delta_{min}}^{\Delta_{max}} \sum_{f=1}^{n_f} J_{ij}^{inh} \int_0^{\infty} ds \varepsilon(s) \delta(t - s - \Delta^{ax} - t_f^j) \quad (8)$$

In contrast to (7), the inhibitory synaptic efficacy J^{inh} is identical for all neurons and all synapses. It is adjusted according to the firing probability p in the spike patterns the system has been trained on. During the simulations of Sect. 3 it is set to $J^{inh} = p = 0.025$.

Equations (7) and (8) show that the postsynaptic neuron sums over the signals from all other neurons (spatial average) and integrates over the past (temporal average). Based on this observation, it has often been argued that quantities that average over space (ensemble activity) or time (mean firing rate) should be sufficient to describe the states of the network. The results of Sect. 3, will show, however, that the network is capable of storing and retrieving patterns with high resolution in space and time – despite the postsynaptic averaging. This clearly underlines the potential relevance of non-averaged or "local" information in neural networks.

In order to store those time-resolved excitation patterns, the efficacies have to be adjusted according to a "Hebbian" learning rule. How this can be done will be explained now.

2.3 Hebbian learning

As a basic principle of synaptic learning, Hebb (1949) postulated that synaptic efficacies are enhanced during




Fig. 3a, b. Hebbian learning at the synapse. The presynaptic neuron j fires at time t_j^f and the postsynaptic neuron i at t_i^f . It takes a time Δ^{ax} and Δ^{dent} , respectively, before the signal arrives at the synapse. At the presynaptic terminal neurotransmitter is released (shaded) and evokes an EPSP (dashed) at the postsynaptic neuron. In **a** the dendritic spike arrives slightly after the neurotransmitter release and matches the time window defined by some chemical processes, so the synaptic efficacy is enhanced. In **b** the postsynaptic neuron has fired too early and no strengthening of the synapse occurs

periods of simultaneous activity of the presynaptic and postsynaptic neurons. Long-term potentiation (LTP), which has been found in hippocampal brain slices, seems to support this postulate (Kelso et al. 1986; Brown et al. 1989, 1991). In most experiments on LTP and in many models of Hebbian learning the terminology of "simultaneous activity" is understood as a time-averaged quantity, e.g., mean firing rate or mean membrane voltage. In our approach we shift emphasis to single spikes and EPSPs and consider the momentary state of pre- and postsynaptic membrane directly at the synapse. The main idea is that only those quantities which are locally available at the synapse can lead to a change of synaptic efficacy (Herz et al. 1988, 1989).

In order to make these ideas more precise, we consider a synapse J_{ij}^{ax} from neuron j to neuron i (see Fig. 3). If neuron j fires at time t_j^f , a spike is transmitted along the axon and arrives after a time Δ^{ax} at the synapse of strength J_{ij}^{ax} . Here it induces the release of neurotransmitter substance which finally evokes an EPSP at the postsynaptic neuron i . Let us now assume that neuron i has also fired, but at a different time t_i^f . When and how does information about the postsynaptic spiking arrive at the synapse? As a solution to this question we tentatively suggest that firing causes not only emission of an action potential along the axon, but also some dendritic

pulse that spreads by (active or passive) transport over the dendritic tree. While active transport is known in the hippocampus (Wong et al. 1979), it has, so far, not been seen in the neocortex. The relevance of these pulses to synaptic coding has not been discussed so far. The important point in this context is that active processes would allow the pulse to remain a point-like event, local in space and time. A synapse could thus receive precise time-resolved information on postsynaptic spiking.

As to synaptic learning, we propose that a strengthening of the synapse occurs only if the dendritic pulse arrives at the synapse concurrently with (or slightly after) the neurotransmitter (or some other messenger substance) released from the presynaptic terminal. At this point our assumption is purely speculative. Most experiments have studied, not the detailed time structure of LTP, but rather mean activity. Dendritic pulses have not yet been considered in this context. Let us assume that the arrival of transmitter substance peaks within 1 ms after the beginning of presynaptic release. This defines a narrow time window $v(t)$ for a match with the dendritic pulse:

$$v(t) = \exp \left[-\frac{1}{2} \left(\frac{t - \Delta^{chem}}{\tau_{chem}} \right)^2 \right] \quad (9)$$

(see also Fig. 3). Synaptic efficacy is enhanced only if the signal of the dendritic pulse arrives during the time of transmitter release $v(t)$ due to a presynaptic spike. This is our precise interpretation of the phrase “simultaneous activity”.

Adaptation of synaptic efficacies is done during a separate learning session of total length $T = \sum_{\mu=1}^q T^\mu$. During the training period T^μ a spatio-temporal excitation pattern μ with $1 \leq \mu \leq q$ is forced upon the network. Let $t_i^f(\mu)$ and $t_j^g(\mu)$ denote the firing times of the post- and presynaptic neuron. The presynaptic spike needs a time Δ^{ax} to travel to the synapse J_{ij}^{ax} , the postsynaptic pulse arrives there after a time Δ^{dent} . To keep arguments simple, we assume that the dendritic delay Δ^{dent} is independent of the position of the synapse on the dendritic tree. In view of the above considerations, a change of efficacy for a synapse with delay Δ^{ax} is – after the presentation of pattern μ – given by

$$\begin{aligned} \delta J_{ij}^{\Delta^{ax}}(\mu) \sim & \sum_{f=1}^{F_i^\mu} \int_0^{T^\mu} dt \delta(t - \Delta^{dent} - t_i^f(\mu)) \\ & \times \sum_{g=1}^{F_j^\mu} \int_0^\infty ds v(s) \delta(t - \Delta^{ax} - t_j^g(\mu) - s) \end{aligned} \quad (10)$$

The total number of spikes neuron i has to fire to generate the pattern μ is denoted by F_i^μ . The first δ function on the right-hand side of (10) describes the postsynaptic activity at the location of the synapse, the second δ function the presynaptic activity. A strengthening of the synapse occurs only if the two activities coincide within a time window defined by $v(s)$. The integrals over time can be performed out, yielding

$$\delta J_{ij}^{\Delta^{ax}}(\mu) = \alpha^\mu \sum_{f=1}^{F_i^\mu} \sum_{g=1}^{F_j^\mu} v(t_i^f(\mu) - t_j^g(\mu) + \Delta^{dent} - \Delta^{ax}) \quad (11)$$

where α^μ is some proportionality factor. It is a free parameter and can be used to normalize the right-hand side of (11). To keep things simple we assume that all neurons have the same mean firing rate and fire exactly $F_i^\mu = F_j^\mu \equiv F^\mu$ spikes during the training session T^μ . A useful normalization is then $\alpha^\mu = 1/F^\mu$.

Several spatio-temporal excitation patterns can be learned subsequently by the same procedure. We assume that the final strength of the synapse is given by a linear superposition of all contributions:

$$\begin{aligned} J_{ij}^{\Delta^{ax}} = \sum_{\mu=1}^q \delta J_{ij}^{\Delta^{ax}}(\mu) = \sum_{\mu=1}^q \sum_{f,g=1}^{F^\mu} \frac{1}{F^\mu} v(t_i^f(\mu) - t_j^g(\mu) \\ + \Delta^{dent} - \Delta^{ax}) \end{aligned} \quad (12)$$

After the training session the efficacies are kept at a fixed value. Retrieval of the excitation patterns can be induced by a short external signal (see Sect. 3).

Before proceeding to simulations of a special system we would like to explain how the system works in the general case. A neuron which has fired at $t = 0$ is inhibited thereafter by its own refractory field. As a consequence it has a reduced excitability (higher threshold). Only if the post-synaptic potential, $h_i^{syn}(t)$ is strong enough to exceed threshold, $h_i^{syn}(t) + h_i^{ref}(t) > \theta$, can it




Fig. 4. Learning and retrieval. Neuron i , which has fired at $t = 0$ (marked by a black dot), is subject thereafter to a refractory field $h_i^{ref}(t)$ equivalent to an increase of a dynamic threshold $\theta - h_i^{ref}(t)$ (dashed). It fires again if the postsynaptic potential is strong enough, viz., $h_i^{syn}(t) = \theta - h_i^{ref}(t)$ (top). $h_i^{syn}(t)$ is induced by the action potentials of neurons which have fired a few milliseconds before. The Hebbian learning rule strengthens only those connections with appropriate delay (thick lines) whereas others remain weak (thin lines). A neuron j which has fired much earlier does not contribute at all

fire again. The postsynaptic potential, $h_i^{syn}(t)$, is the sum of all incoming EPSPs. Let us now assume that the network has to learn a spike pattern where neuron i should fire again at t_i^f . The learning rule (12) enhances only those connections where the EPSP has the correct timing so as to cause i to fire at t_i^f (see Fig. 4). This argument shows that the above learning rule should allow the storage of fairly arbitrary spike patterns. For the sake of simplicity we have restricted our simulations to the special case of cyclic patterns. The main results, however, are much more general.

3 Simulation results

3.1 Low-activity random patterns

A network of $N = 1000$ neurons has been trained on $q = 4$ patterns. The procedure easily takes care of many more patterns, but $q = 4$ is convenient for simulation. Each pattern consists of a sequence of spikes from different neurons during a time window of $T = 40$ time steps, i.e., 40 ms. The sequence is then repeated. A spike pattern μ ($1 \leq \mu \leq q$) is defined by exactly one firing time $t_i^f(\mu)$ for each single neuron $1 \leq i \leq N$. The firing time is drawn from a random distribution with equal probability $p = 0.025$ for all times $1 \leq t_i^f(\mu) \leq 40$. Thus, in an ideal and noiseless pattern all neurons fire regularly with a rate of 25 Hz, but the firing of different neurons is randomly correlated.

During the training session all spike patterns $1 \leq \mu \leq q$ are forced upon the neurons and the synaptic efficacies are adjusted according to (12). In order to check whether the patterns are now stable attractors of the neuronal dynamics, retrieval of the patterns has to be studied. A retrieval session is started by a short external stimulus of duration $t_{init} = 5$ ms. It consists of a spatio-temporal sequence of short pulses $h_i^{ext}(t) = \delta[t, t_i^f(\mu) - T]$ for $-5 \leq t < 0$ ms. In other words, the network is

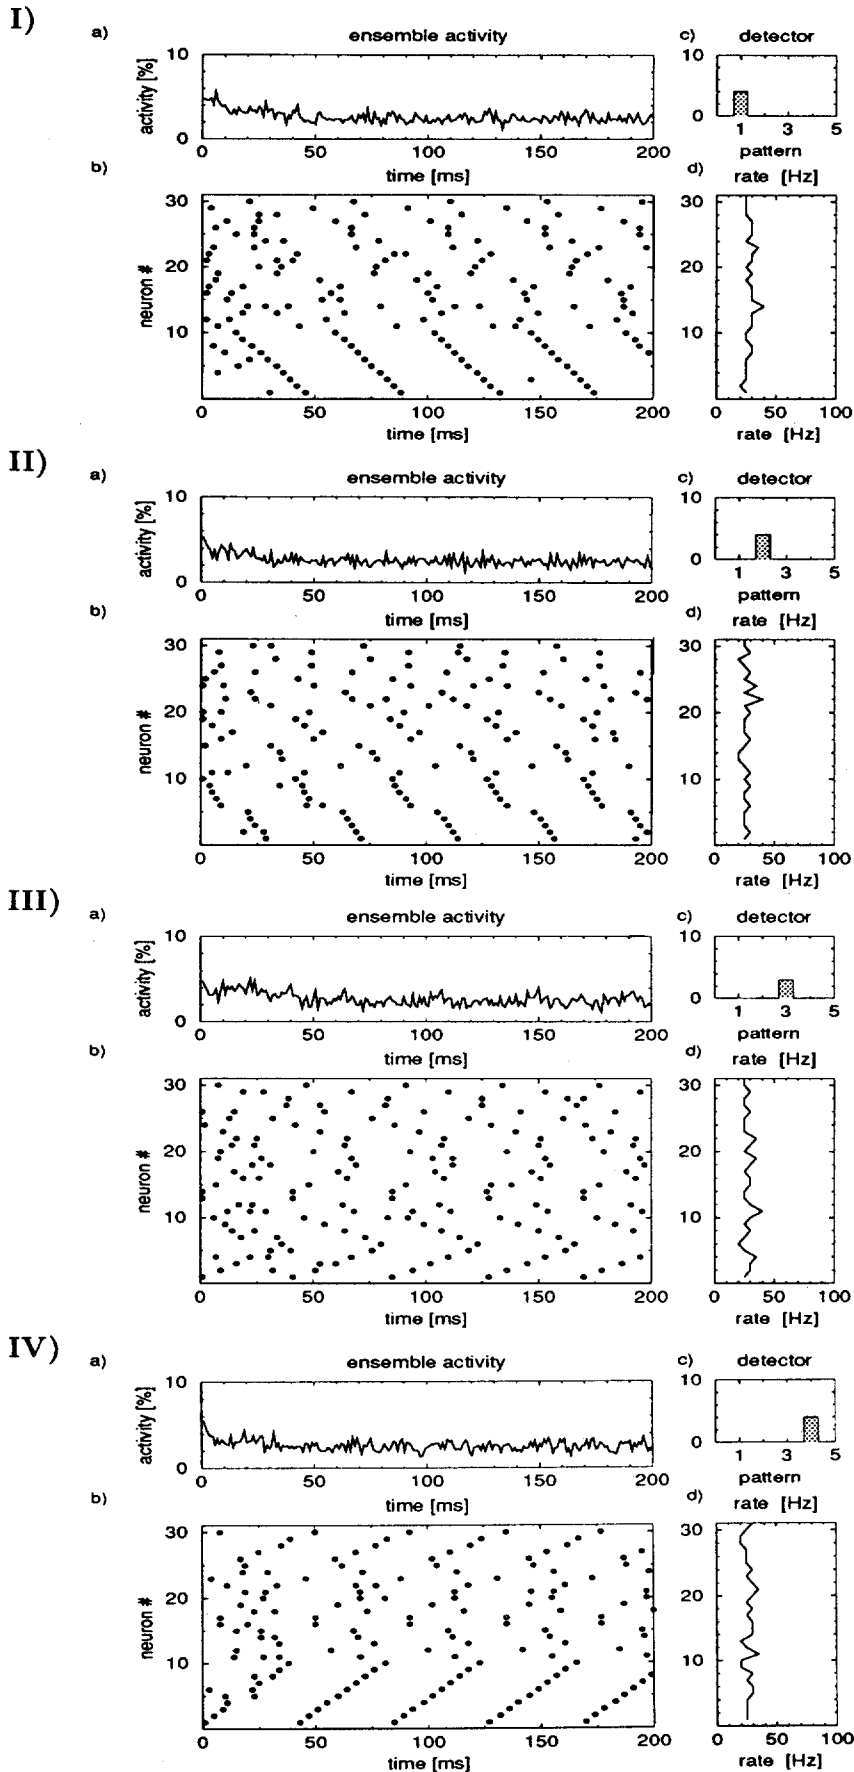


Fig. 5. Retrieval of spatio-temporal spike patterns. Four runs have been plotted, showing the retrieval of four different patterns (*I-IV*). A short (5 ms) simulating pulse just before the start of the recording ($t = 0$) induces the network response shown in *I-IV*. In *a* the ensemble activity averaged over all neurons of the net is plotted. Part *b* shows the spike raster of 30 neurons. Each spike t_i of neuron i (y -axis) is marked as a filled circle at the time of firing (x -axis). For ease of visualization, the lower ten neurons have learned different stripe patterns. Based on the spatio-temporal correlations, the pattern detector *c* classifies the spike pattern. In *d* the mean firing rate (x -axis) of the 30 neurons (y -axis) is shown. Note that both the mean firing rate and the ensemble activity are similar in all runs and are not useful in distinguishing the patterns

initialized during 5 ms in a state consistent with one of the learned patterns. The pattern μ that is matched should be completed and cyclically retrieved afterwards.

The results of four retrieval sessions with different stimuli are shown in Fig. 5. For the four patterns (I–IV), the ensemble activity (spatial average) during retrieval is plotted in panel *a* and the spike raster of the retrieval session is shown in *b*. In *c* the spike pattern is classified, and in *d* the mean firing rate of the neurons is calculated.

Let us discuss the spike raster (Fig. 5Ib–IVb) first. In a way similar to experimental multi-electrode recordings, we have selected 30 neurons which are numbered and plotted along the y -axis. Time is plotted along the x -axis. The origin $t = 0$ marks the end of the stimulus, thus for all $t \geq 0$ we have $h_i^{ext}(t) = 0$. In this subsection, we have also assumed that $h_i^{back}(t) = 0$. All spikes of neuron i appear as black dots along a line parallel to the x -axis. If we count the number of spikes along this line and divide by the total recording time ($T = 200$ ms), we find the mean firing rate of this neuron. This is the quantity which is plotted in *d*, to the right of the spike raster. We see that all neurons have approximately the same firing rate, 25 Hz. Thus, if we consider the mean firing rate only, we can detect no significant structure in the firing behavior of the neurons.

Instead of averaging over time we can also average over space so as to define the ensemble activity of the network. If we count the number of spikes in every millisecond (along a vertical line in the spike raster) and divide by the total number of neurons, we find the ensemble activity plotted in *a*. We see that immediately after the stimulus the ensemble activity is high ($\approx 5\%$), but due to adaptation it decreases. After 50 ms it has settled to an average of 2.5% and no significant structure is left. Nevertheless, if we look at the spike raster (*b*), we see that the network remains in a regular firing state. The specific spike pattern has been induced by the stimulus and is different for parts I–IV of Fig. 5.

For ease of visualization of this fact, we have used a little trick. Neurons with index 1–10 did not learn random patterns but “meaningful” objects such as diagonal stripes (see the spike raster Ib–IVb). Ten neurons hardly disturb a network of $N = 1000$ neurons, but they can help the human reader to find regularities and to recognize different patterns.

The pattern detector (*c*) distinguishes the patterns based on the correlations between the spikes of all neurons in the net. A pattern μ is detected if the firing times during the last 40 ms match the time shifted spike raster of pattern μ . To be more specific, we define a correlation function

$$corr^\mu(t) = \frac{1}{N} \sum_{i=1}^N \sum_{\tau=0}^{39} \delta[t - t_i^f, \tau] \delta[40 - t_i^f(\mu), \tau] \quad (13)$$

where $\delta[x, y]$ denotes the Kronecker δ with $\delta[x, y] = 1$ for $x = y$ and 0 otherwise. A detector signal $D^\mu = +1$ is given if $corr(t) \geq 0.5$ during the simulation run ($0 < t < 200$). Thus we require that 50% of the spikes must be correlated exactly with spike pattern μ .

If we analyze the series of Fig. 5, I–IV, a number of conclusions can be drawn. First of all, it is indeed possible to store and retrieve spatio-temporal spike patterns with a time resolution of 1 ms in a neural network with biological constraints. This is quite remarkable in view of the typical duration of an EPSP (approximately 5–15 ms). Second, several patterns can be stored in the same network. In analogy to the storage capacity calculations of Gardner (1988) we expect that the number of spatio-temporal spike patterns of duration T that can be stored in an optimally designed network of N neurons is proportional to $NC/(Tp|\ln p|)$, where C is the connectivity per neuron (here $C = 4$) and $p \ll 1$ is the globally averaged activity of a typical spatio-temporal pattern. From the static case it is known that optimal storage requires a long iterative learning procedure (Krauth and Mézard 1987), whereas a standard single-step learning rule yields a reduced capacity $\alpha_c^H = 0.14$ (Amit et al. 1987). The same result, $\alpha_c = 0.14$, also holds for cyclic patterns in a model network with a given distribution of delays (Herz et al. 1991). These arguments lead to the estimation that our model network of $N = 1000$ neurons with $p = 1/T = 0.025$ should be capable of storing roughly 100 patterns. Exact answers to the question of storage capacity are presently under investigation.

Comparing panels *d* and *a* in Fig. 5, I–IV, we see that neither the mean firing rate nor the ensemble activity contains significant structure that would suffice to distinguish the four patterns. Nevertheless, the time-resolved spike raster shows that the four patterns can be retrieved and distinguished. This is confirmed by the pattern detector, which measures the correlations in the spiking of different neurons, but, of course, this could also be done by a neuron or cluster of neurons which has learned exactly this classification in the same way as described before. Methods that are based on mean firing rates or ensemble activities would miss the information contained in the time-resolved spike raster. The above examples clearly show that single spikes can carry important information.

3.2 PSTH analysis

The arguments used so far have been based on the results of a single run of the experiment. If several runs are made, the resulting spike raster may, and usually will, look different every time. This is partly due to noise in neurons and synapses that induces variability on a fast time scale. Much greater variability, however, comes from the fact that other areas of the neural system which are out of control of the experimenter influence the network state. In real systems, only a small portion of the brain is connected directly to the sensory receptors. The major part is involved with further processing of incoming signals and acts somewhere in the background. Due to feedback connections, the processing state of the “background” influences the firing state of the primary areas.

We have modelled this effect very crudely by a single feedback signal $h_i^{back}(t)$ that acts identically on all neurons in the “primary” network and varies only slowly in time. To be specific, we now assume that the background signal $h_i^{back}(t)$ varies stochastically every 120 ms




Fig. 6. Post-stimulus time histogram (PSTH). An identical stimulus was applied in 15 runs of the experiment. The stimulating pulse ended just before the start of the recording ($t = 0$). The input of the background channel varied every 120 ms. The response of neurons 1 and 8 in 15 repetitions has been plotted as a PSTH (binwidth 5 ms). The PSTH shows a structured response only during 600 ms. Afterwards the signal is buried in noise; cf. Fig. 7, which clearly contradicts the conclusion that all the information on the stimulus is lost

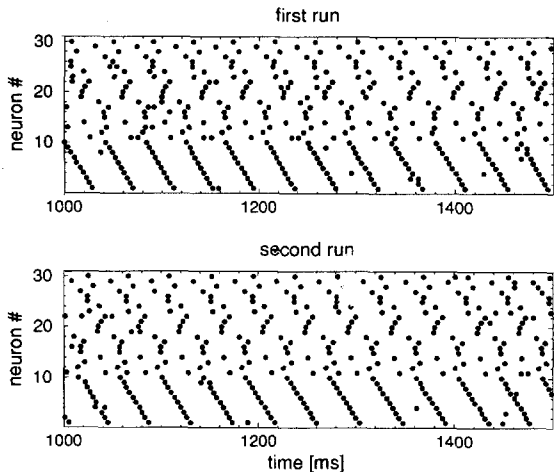


Fig. 7. Spike raster during two of the 15 runs of Fig. 6. The spikes of 30 neurons are plotted between 1000 and 1500 ms after the stimulating pulse. Both spike rasters exhibit the diagonal stripes of pattern 1; cf. Fig. 6. The spike raster of a single run carries information which is hidden in the noisy part of the PSTH of Fig. 6 ($1000 \leq t \leq 1500$)

in a range of ± 0.15 i.u. with equal probability [$h_i^{back}(t)$ was assumed to be zero in Sect. 3.1]. This rather long variation time has been chosen in order to distinguish the background signal from the effects of noise in each neurons; see (1). An 8–10 Hz input also looks reasonable in view of a possible thalamic origin of these variations.

The results of 15 runs over 1500 ms are added and a PSTH is computed for neurons 1 and 8 (Fig. 6). For both neurons we find a clear oscillatory (stimulus locked) structure with a frequency about 20 Hz, which is the

same as in Fig. 5. It is sustained during the first 600 ms. For times longer than 800 ms the signal disappears in a noisy background.

It is now tempting to conclude that the network which has been triggered by an external stimulus at $t = 0$ has lost all information about the stimulus after 800 ms. This conclusion, however, is wrong, as is shown in Fig. 7, where the spike raster in two of the 15 runs is plotted for times between 1000 and 1500 ms. We clearly see that the spike pattern is still fully sustained – despite the fact that the PSTH contains no information for $t \geq 800$ ms. An explanation of this difference can be found if we compare the spike rasters of the first and second run of Fig. 7 in more detail. There is a phase shift in the spike pattern of the two runs. Adding several spike patterns with different phases yields the “noisy” PSTH of Fig. 6.

To summarize this subsection, we have shown in a simple example that information which is contained in a single run of the experiment may be lost in the PSTH which is accumulated over several runs. Due to background activity which cannot be controlled by the experimenter, the response to a stimulus is usually different from one run to the next. Averaging over several runs may then spoil the results. Because of this we would like to emphasize the importance of the spike raster of a single run.

3.3 Rate quantization

In which way does the signal from the background affect the spike pattern? To answer this question we now adopt an approach that allows us – in contrast to the experimental situation – to control the background activity. Indeed, in a computer experiment we can change the background activity $h_i^{back}(t)$ systematically from low to high. The results have been plotted in Fig. 8.

The external stimulus is, as before, a short 5 ms pulse and is identical in each run. In the first run, Fig. 8I, background activity is low, $h_i^{back}(t) \equiv 0.05$ i.u. The result is the same as in Fig. 5a, except for a phase shift of the response. If the background activity is high, $h_i^{back}(t) \equiv 0.2$ i.u., the network is driven into a state with *doubled* firing rate and twice the activity (Fig. 8, II). If the background input is even stronger, $h_i^{back}(t) \equiv 0.35$ i.u., rates can even be three times as high (Fig. 8III). We call this phenomenon *rate quantization* because a pattern is recalled n times as often as learned where n is a positive integer ($n \geq 0$). Nevertheless, a glance at the spike raster b reveals that the pattern can still be recognized as a spatio-temporal entity. Figure 8, II is the same as Fig. 8, I except that the pattern evolves twice as often. This is confirmed by the pattern detector c , which counts twice as many instances of pattern 1. In Fig. 8, III, the pattern detector fails, but a closer look at the spike raster reveals that pattern 1 (“diagonal stripes for the lower 10 neurons”) can still be identified. It now runs three times as often in a given time interval of, say, 40 ms. Keeping its time structure, the pattern appears once (I), twice (II), or thrice (III) during this time interval. In short, it does not run faster but, instead, it appears more often.

In our model this type of rate quantization is due to the fact that the prototype of each spatio-temporal

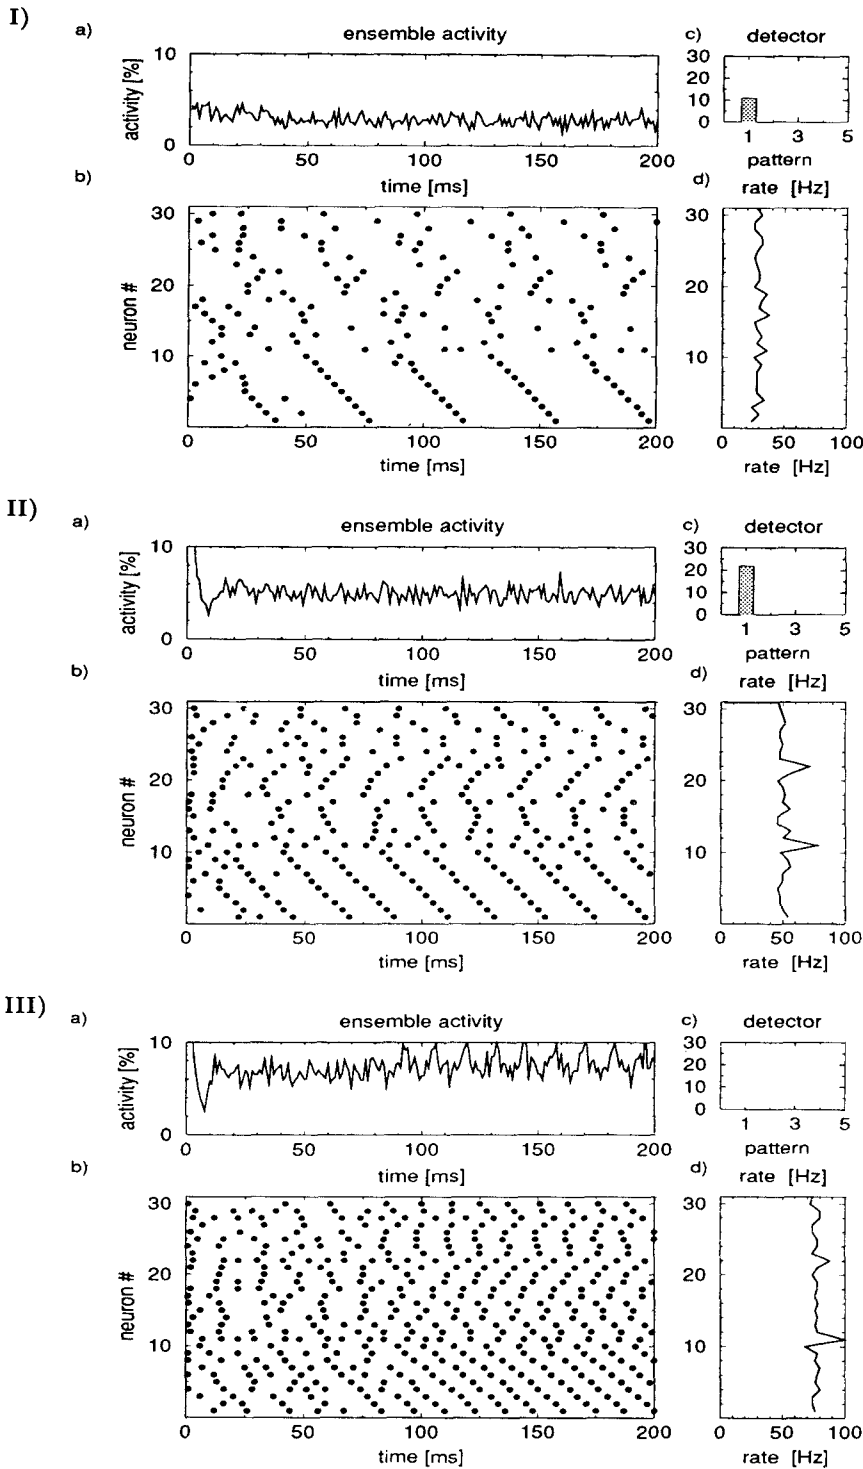


Fig. 8. Rate quantization. Frequency doubling and tripling appears under the influence of a background signal, say, from a higher cortical area. *a* Ensemble activity, *b* spike raster, *c* pattern detector, and *d* mean firing rate are as in Fig. 5. A 5-ms pulse stimulating pattern 1 is applied just before the start of the recording. If the level of the background activity is changed, mean firing rate and ensemble activity can be twice (*II*, $h_i^{back}(t) = 0.20$ i.u.) or even three times (*III*, $h_i^{back}(t) = 0.35$ i.u.) as high as in the reference run (*I*, $h_i^{back}(t) = 0.05$ i.u.). Nevertheless the time structure of the excitation pattern (diagonal stripes matching pattern 1) is always the same

pattern consists in a closed cycle of length T , i.e., the state at time $t + T$ is identical to the state at t . Higher rates correspond to the second and third harmonics of the basic cycle. In real biological systems it is rather unlikely that a spatio-temporal spike pattern is a closed cycle of definite length. Instead it could consist in a long spatio-temporal sequence that finally stops. It is nevertheless

tempting to speculate that a similar rate quantization exists. This could be the case if, for high background activity, each spike at time $t_i^f(\mu)$ is replaced by a short burst of two or three spikes while the overall pattern is preserved. To phrase it differently, the coarse time structure of the spike pattern may be always the same, despite large differences in the rates.

4 Discussion

We have presented a rather crude model of a neural network with time-resolved learning and retrieval capabilities. Spike emission and reception is described by two response functions, viz., the refractory function $\eta(s)$ and the EPSP $\varepsilon(s)$ (spike response model, SRM). In our present network, there is only a single type of model neuron, whereas several classes of neurons with a variety of intrinsic characteristics are found in real neural structures. Each pair of neurons in our model is connected by exactly four synapses, whereas the connectivity in real neural systems is much more complex. Inhibition has been modelled in a coarse way by a single activity-dependent neuron, whereas in reality inhibition is local. We have described a homogeneous network, while natural brains are organized in several distinct areas and layers. Areas in the background have been mimicked by a single feedback channel which sends identical signal to all neurons in the "primary" network. In all these respects, our model network can only be a rough approximation of real neural systems.

On the other hand, our modelling approach is carefully crafted with regard to the time constants of neural systems. We have assumed realistic axonal delays ranging from 1 to 4 ms. We have also attempted to describe the EPSP realistically (see Fig. 2). Our model EPSP has a rise time of 3 ms and a slowly decaying tail. The overall duration (width of half-maximum) is approximately 8 ms, in good agreement with experimental data (Brown et al. 1989). The model of a single neuron in our network includes an absolute refractory time of 3 ms and an AHP potential with long time constants. The exact numerical values are somewhat arbitrary, but they lie in a reasonable regime. As a result, both spike train and gain function have a realistic shape (Gerstner and van Hemmen 1992a).

Based on our careful modelling of biological time constants, we have tried to answer the central question of this paper: What is the physiologically relevant time scale of neuronal signal processing? Regarding the long postsynaptic integration time and spatial summation performed on the dendritic tree, it has sometimes been argued that relevant changes should occur on a time scale of 100 ms and more (Amit and Tsodyks 1991). The results of Sect. 3, however, show that it is possible to store and retrieve information with a time resolution of 1 ms.

Description of signal transmission times also allows some insight into the timing necessary for Hebbian learning of spike patterns. Our Eq. (12) requires that precise information on the firing of pre- and postsynaptic neurons be available at the synapse. A similar formula for Hebbian learning of spatio-temporal excitation patterns has been derived by Herz et al. (1988, 1989). We note, however, two major differences. First, Herz et al. assume a much slower time scale for the overall performance of their network. Temporal changes of the activity pattern occur during times which are long compared to the dendritic integration time and internal time constants of the neurons. To induce temporal sequences they consequently need very long axonal, or other, delays (up to

20 ms), a fact which has been subject to criticism. The present model, however, uses short axonal delays (between 1 and 4 ms), for which there is stronger experimental evidence. The second difference stems from the distinction between electrical and chemical processes which we make in the present model. This allows us to apply a narrow coincidence window – described by the messenger release $v(\tau)$ – to define simultaneous pre- and postsynaptic activity. If LTP were induced by electrical processes only, the electrical response $e(\tau)$ would define the time window – acting on a much slower time scale. The chemical processes, however, which somehow induce the electrical response may be much sharper than the EPSP. Since the sharp presynaptic signal is combined with an equally sharp dendritic pulse we get the high time resolution necessary to store fast spatio-temporal excitation patterns.

By way of example we have shown that retrieval of spike patterns with a resolution of 1 ms is possible despite the much longer dendritic integration time. This is a theoretical result based mainly on computer simulations. It does not imply that all biological systems actually work on such fast a time scale. On the other hand, we may ask why nature should not make use of this possibility. Experience tells us that evolution grasps most opportunities of improvement. In all cases where a "fast" system is somehow "better" than a slow one, we would therefore expect that single spikes and their cross-correlations become important. A fly rushing around at high speed through an unknown environment is an example of an animal that needs a fast reaction time. Indeed, it has been shown experimentally that spikes are of eminent importance for signal processing in the fly (Bialek et al. 1991).

If it is true that single spikes are relevant, some traditional methods of data analysis and several assumptions of model approaches need to be reconsidered carefully. Future experiments should be checked for the appropriate time resolution. This can be done by analyzing the information content of the experimental data with variable resolution. A time resolution of down to 1 ms could be necessary, but a resolution of 10 or 30 ms seems to be enough in most cases where such an analysis has been carried through so far (Optican and Richmond 1987; Krüger and Becker 1991; Eskandar et al. 1992). Post-stimulus time histograms should be used for time-resolved data analysis only if the sources of noise are known and if responses do not vary too much from trial to trial.

Similar considerations hold for theoretical approaches to network modelling. In vertebrates, a network of analog neurons should be used only in those cases where the relevant states can be defined in terms of mean firing rates. The standard model of associative memory (Hopfield 1982) is an example of a system whose stationary behavior is independent of the model neuron (Hopfield 1984; Gerstner and van Hemmen 1992b). In general, stationary states can be fully described by mean firing rates. In all time-dependent problems, however, spiking is essential and the type of model neuron that is used as the basic unit of a neural network has to be chosen

carefully. In particular, if activity patterns that change on a time scale of 25 ms or less are considered (as in the case of collective oscillations in the cortex), the time constants of spiking and signal transmission become important (Gerstner and van Hemmen 1992a; van Hemmen et al. 1992; Gerstner et al. 1993). In these cases, a spiking model neuron should be preferred.

It is, of course, a trivial statement that information is lost during averaging. With regard to the methods frequently applied in experimental and theoretical neurobiology, however, we felt a need to discuss the implications of this abstract statement. The above examples are intended to give a deeper understanding of the problems with traditional techniques of data analysis and neural modelling. Mean firing rate, ensemble activity, and PSTH analysis all involve some kind of averaging and should therefore be applied only if care has been taken regarding the consequences. To phrase it succinctly: Beware of a priori averaging – you may miss the essentials.

Acknowledgments. This work has been supported by the Deutsche Forschungsgemeinschaft (DFG) under Grant No. He 1729/2-1. Raphael Ritz holds a scholarship of the Freistaat Bayern.

References

- Abeles M (1982) Local cortical circuits. Springer, Berlin Heidelberg New York
- Abeles M, Lass Y (1975) Transmission of information by the axon. *Biol Cybern* 19:121–125
- Adrian ED (1926) The impulses produced by sensory nerve endings. *J Physiol (Lond)* 61:49–72
- Aertsen A, Gerstein G, Johannesma P (1986) From neuron to assembly: neuronal organization and stimulus representation. In: Palm G, Aertsen A (eds) *Brain theory*. Springer, Berlin Heidelberg New York pp 7–24
- Amit DJ, Tsodyks MV (1991) Quantitative study of attractor neural networks retrieving at low spike rates. I. Substrate spike rates and neuronal gain. *Neural Computation* 2:259–273
- Amit DJ, Gutfreund H, Sompolinsky H (1985) Spin-glass models of neural networks. *Phys Rev A* 32:1007–1032
- Amit DJ, Gutfreund H, Sompolinsky H (1987) Statistical mechanics of neural networks near saturation. *Ann Phys (NY)* 173:30–67
- Bialek W, Rieke F, Ruyter van Steveninck RR de, Warland D (1991) Reading a neural code. *Science* 252:1854–1857
- Brown TH, Ganong AH, Kairiss EW, Keenan CL, Kelso SR (1989) Long-term potentiation in two synaptic systems of the hippocampal brain slice. In: Byrne JH, Berry WO (eds) *Neural models of plasticity*. Academic Press, San Diego, pp 266–306
- Brown TH, Zador AM, Mainen ZF, Claiborne BJ (1991) Hebbian modifications in hippocampal neurons. In: Baudry M, Davis JL (eds) *Long-term potentiation*. MIT Press, Cambridge, Mass, pp 357–389
- Creutzfeldt OD (1983) *Cortex cerebri*. Springer, Berlin Heidelberg New York, pp 76–78
- Eckhorn R, Grüsser OJ, Kröllner J, Pellnitz K, Pöpel B (1976) Efficiency of different neural codes: information transfer calculations for three different neural systems. *Biol Cybern* 22:49–60
- Eckhorn R, Bauer R, Jordan W, Brosch M, Kruse W, Munk M, Reitboeck HJ (1988) Coherent oscillations: a mechanism of feature linking in the visual cortex? Multiple electrode and correlation analyses in the cat. *Biol Cybern* 60:121–130
- Eskandar EN, Richmond BJ, Hertz JA, Optican LM, Troels K (1992) Decoding of neuronal signals in visual pattern recognition. (*Advances in neural information processing* 4) Morgan Kaufmann, San Mateo, Calif
- Gardner E (1988) The space of interactions in neural network models. *J Phys A: Math Gen* 21:257–270
- Gerstein GL, Perkel DH (1972) Mutual temporal relations among neuronal spike trains. *Biophys J* 12:453–473
- Gerstein GL, Bloom MJ, Espinosa IE, Evanczuk S, Turner MR (1983) Design of a laboratory for multineuron studies. *IEEE Trans Syst Man Cybern* 13:668–676
- Gerstner W (1993) *Kodierung und Signalübertragung in neuronalen Systemen: Assoziative Netzwerke mit stochastisch feuernenden Neuronen* Thesis, Nov 1992, TU Munich. Harri Deutsch, Frankfurt
- Gerstner W, Hemmen JL van (1992a) Associative memory in a network of 'spiking' neurons. *Network* 3:139–164
- Gerstner W, Hemmen JL van (1992b) Universality in neural networks: the importance of the mean firing rate. *Biol Cybern* 67:195–205
- Gerstner W, Ritz R, Hemmen JL van (1993) A biologically motivated and analytically soluble model of collective oscillations in the cortex. I. Theory of weak locking. *Biol Cybern* 68:363–374
- Hebb DO (1949) *The organization of behavior*. Wiley, New York
- Hemmen JL van, Gerstner W, Ritz R (1992) A 'microscopic' model of collective oscillations in the cortex. In: Taylor JG, Caianiello EK, Cotterill RNJ, Clark JW (eds) *Neural network dynamics*. Springer, Berlin Heidelberg New York, pp 250–257
- Herz A, Sulzer B, Kühn R, Hemmen JL van (1988) The Hebb rule: representation of static and dynamic objects in neural nets. *Europhys Lett* 7:663–669
- Herz A, Sulzer B, Kühn R, Hemmen JL van (1989) Hebbian learning reconsidered: representation of static and dynamic objects in associative neural nets. *Biol Cybern* 60:457–467
- Herz AV, Li Z, Hemmen JL van (1991) Statistical mechanics of temporal association in neural networks with transmission delays. *Phys Rev Lett* 66:1370–1373
- Hodgkin AL (1948) The local electric changes associated with repetitive action in a non-medullated axon. *J Physiol (Lond)* 107:165–181
- Hopfield JJ (1982) Neural networks and physical systems with emergent collective computational abilities. *Proc Natl Acad Sci USA* 79:2554–2558
- Hopfield JJ (1984) Neurons with graded response have computational properties like those of two-state neurons. *Proc Natl Acad Sci USA* 81:3088–3092
- Hubel DH, Wiesel TN (1962) Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J Neurophysiol* 28:215–243
- Jack JJB, Noble D, Tsien RW (1975) *Electric current flow in excitable cells*. Clarendon Press, Oxford
- Johannesma P, Aertsen A, Boogaard H van den, Eggermont J, Epping W (1986) From synchrony to harmony: ideas on the function of neural assemblies and the interpretation of neural synchrony. In: Palm G, Aertsen A (eds) *Brain theory*. Springer, Berlin Heidelberg New York, pp 25–48
- Kelso SR, Ganong AH, Brown TH (1986) Hebbian synapses in hippocampus. *Proc Natl Acad Sci USA* 83:5326–5330
- Krauth W, Mézard M (1987) Learning algorithms with optimal stability in neural networks. *J Phys A: Math Gen* 20:L745–L752
- Krüger J (1983) Simultaneous individual recordings from many cerebral neurons: techniques and results. *Rev Physiol Biochem Pharmacol* 98:177–233
- Krüger J, Aiple F (1988) Multimicroelectrode investigation of monkey striate cortex: spike train correlations in the infragranular layers. *J Neurophysiol* 60:798–828
- Krüger J, Becker JD (1991) Recognizing the visual stimulus from neuronal discharges. *Trends Neurosci* 14:282–286
- Kühn R, Bös S, Hemmen JL van (1991) Statistical mechanics for networks of graded-response neurons. *Phys Rev A* 43:2084–2087
- Lancaster B, Adams PR (1986) Calcium-dependent current generating the afterhyperpolarization of hippocampal neurons. *J Neurophysiol* 55:1268–1282
- MacKay DM, McCulloch WS (1952) The limiting information capacity of a neuronal link. *Bull Math Biophys* 14:127–135
- Miyashita Y (1988) Neuronal correlate of visual associative long-term memory in the primate temporal cortex. *Nature* 335:817–820
- Optican LM, Richmond BJ (1987) Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. 3. Information theoretic analysis. *J Neurophysiol* 57:162–178

- Palm G, Aertsen AMHJ, Gerstein GL (1988) On the significance of correlations among neuronal spike trains. *Biol Cybern* 59:1-11
- Perkel DH, Gerstein GL, Moore GP (1967) Neuronal spike trains and stochastic point processes. I. The single spike train. II. Simultaneous spike trains. *Biophys J* 7:391-418, 419-440
- Reitboeck HJA (1983) A multi-electrode matrix for studies of temporal signal correlations within neural assemblies. In: Basar E, Flohr H, Haken H, Mandell AJ (eds) *Synergetics of the brain*. (Springer series in synergetics 23) Springer, Berlin Heidelberg New York, pp 174-182
- Rumelhart DE, McClelland GL (eds) (1986) *Parallel distributed processing, vol 1: Foundations*. MIT Press, Cambridge, Mass
- Stein RB (1967) The information capacity of nerve cells using a frequency code. *Biophys J* 7:797-826
- Wong RKS, Prince DA, Basbaum AF (1979) Intradendritic recordings from hippocampal neurons. *Proc Natl Acad Sci USA* 76:986-990