Proceedings of the 3rd Joint Sympostuim on Neural Computation, La Jolla, CA
Vol. 6, 161-171 (1996)

Correlation Coding in a Stochastic Network Model of Auditory Binding

Raphael Ritz and Terrence J. Sejnowski
ritz@salk.edu and terry@salk.edu
Computational Neurobiology Laboratory
The Salk Institute for Biological Studies
10010 North Torrey Pines Road, La Jolla, CA 92037, USA

Abstract

DeCharms et al. (1995) have provided evidence for stimulus-dependent changes in the correlations between spike trains of simultaneously-recorded pairs of neurons from the auditory cortex of marmosets even when there was no change in the average firing rates. Most of the characteristics of these experimental observations can be reproduced by a simple model based on neurons having leaky integration, fire-and-reset spikes and with Poisson-distributed, balanced input. The source of synchrony in the model was common sensory input. Spike frequency adaptation was implemented by sensory-driven, delayed inhibition. The outputs of neurons in the model appear noisy (almost Poisson) owing to the stochastic nature of the input signal, but there is nevertheless a strong central peak in the correlation of the output spike trains. The experimental data and this simple model clearly demonstrate how even a noisy-looking spike train can convey basic information about a sensory stimulus in the relative spike timing between neurons. We address the binding problem and show why synchrony without periodicity might be advantageous in representing multiple objects at the same cortical site simultaneously.

1 INTRODUCTION

It is commonly believed that the neural code used by nerve cells to transmit information in the cerebral cortex is the mean firing rate of action potentials. Whereas there is solid evidence for this coding scheme at the neuromuscular junction, where this concept originated, the temporal averaging involved in the decoding process causes problems at the cortical level, where neurons usually fire at rates too low to allow for a sufficiently long decoding

time. As a possible solution to this problem it has been proposed that cells could also perform a spatial average instead of, or in addition to, temporal averaging. But this form of population code also assumes that information is coded in a firing *rate*—whether spatial or temporal—and that a neuron simply reflects changes in its input firing rates by modulating its output firing rate. This is the underlying assumption allowing the common reduction to a transfer function used by most artificial neural network models to describe single neuron processing.

Recently, deCharms et al. (1995) presented evidence for a different form of coding in the primary auditory cortex of marmosets. They showed that rapidly adapting cells responded to elongated tone stimuli with a fast transient onset response returning quickly to spontaneous firing rates. Thus, these cells cannot convey information about a steady-state stimulus by their firing rate. However, these cells do show an increase in their tendency to fire *simultaneously* as revealed by correlation analysis if they are tuned to the presented stimulus frequency. Nevertheless, each spike train looked almost like it was randomly generated and there was no stimulus-locked component as shown by a flat shift predictor.

Most characteristics of these experimental findings can be reproduced in a simple neuronal model using leaky integrate-and-fire units with Poisson-distributed, balanced input, as shown below.

2 THE RANDOM WALK MODEL

Assume that the generation of action potentials relies on the membrane potential $u_i(t)$ of cell i ($1 \leq i \leq N$) at time t crossing a firing threshold θ and that deviations from the resting potential (set to 0 here) are due to an input current $C_i(t)$ and given these deviations decay exponentially with the membrane time constant τ_m. The following equation governs the temporal evolution of the membrane potential:

$$\frac{\mathrm{d}}{\mathrm{d}t}u_i(t) = -\frac{1}{\tau_m}u_i(t) + C_i(t) \ . \tag{1}$$

A spike occurs when $u_i(t) = \theta$, and u_i is reset to its resting level. To avoid unrealistically large hyperpolarizations, we also introduce a negative saturation limit θ^{inh}, i.e., we assure $u_i(t) \geq \theta^{\mathrm{inh}}$ for all t. To specify the input current, $C_i(t)$, assume that this input can be subdivided into a background and a stimulus component, $C_i^{\mathrm{bg}}(t)$ and $C_i^{\mathrm{stim}}(t)$ respectively

$$C_i(t) = C_i^{\mathrm{bg}}(t) + C_i^{\mathrm{stim}}(t) \ , \tag{2}$$

and that each of these components consists of excitatory as well as inhibitory parts

$$C_i^{\mathrm{bg,stim}}(t) = E_i^{\mathrm{bg,stim}}(t) - b^{\mathrm{bg,stim}} I_i^{\mathrm{bg,stim}}(t - \Delta^{\mathrm{inh}}) \ , \tag{3}$$

where b denotes a balancing factor indicating the relative strength of the inhibition with respect to the excitation and Δ^{inh} represents a delay. This feedforward network is illustrated for two cells in Fig. 1.

To introduce noise in the model, assume that all excitatory and inhibitory signal components are realizations of an ideal Poisson process, i.e.,

$$E_i^{\mathrm{bg}}(t) = k \ \text{ with probability } \ p(k) = \frac{\lambda^k}{k!}e^{-\lambda} \tag{4}$$

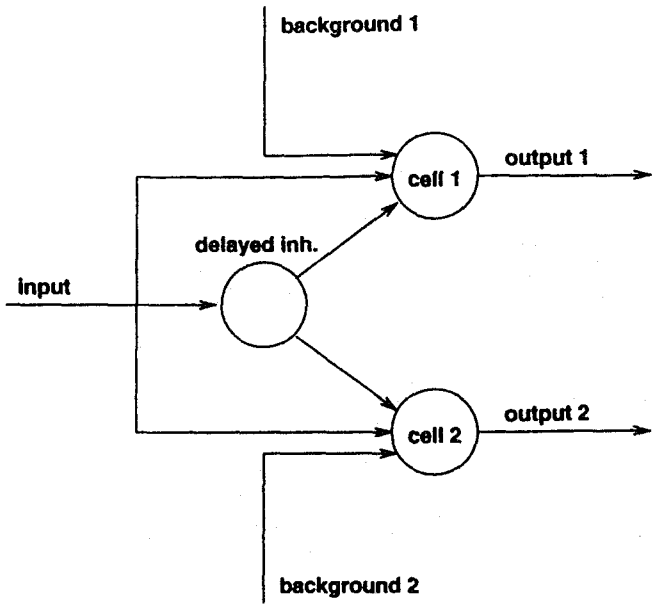

Figure 1: Basic structure for the two cell network. Both cells are getting independent background signals, while the input is assumed to be identical to the two neurons. Both, the input and the background signal, consist of an excitatory and an inhibitory part. The inhibition due to the external stimulus is supposed to arrive later than the excitatory stimulation.

where k is drawn at each time step for every component independently. The parameter λ denotes the mean and the variance of the distribution. Here, it can be interpreted as $\lambda = n_{\text{aff}} \cdot p_{\text{f}}$ the product of the number of afferents times the probability of firing in a single time step, thus fixing the input firing rate. For $\lambda = 10$ and a basic time step of 1 ms, this might correspond to 100 afferents each firing at a rate of 100 Hz.

Due to the randomness in the input the membrane potential undergoes a sort of a random walk with renewal (Gerstein and Mandelbrot 1964).

3 SIMULATION RESULTS

For simulations, we used the following set of parameters. The thresholds were set to $\theta = 15$ and $\theta^{\text{inh}} = -30$, both in units of single EPSP amplitudes. The time scale was fixed by $\tau_{\text{m}} = 10\,\text{ms}$ and $\Delta^{\text{inh}} = 20\,\text{ms}$. We solved (1) using a simple forward Euler method with a time step size of 1 ms. The input was specified by $\lambda = 10$ for all four components and the balancing factors are $b^{\text{bg}} = 1$ and $b^{\text{stim}} = 1.1$. A typical simulation run lasted for three seconds where a stimulus was switched on after the first second and turned off after the second.

3.1 SINGLE CELL PROPERTIES

Consider first the firing of a single cell. As seen in the top row of Fig. 2, the total input current fluctuates vigorously. The resulting membrane potential (Fig. 2, middle) is smoother due to the temporal integration. Threshold crossings of the membrane potential resulting in spike emission were only driven by fluctuations except for the stimulus onset period,

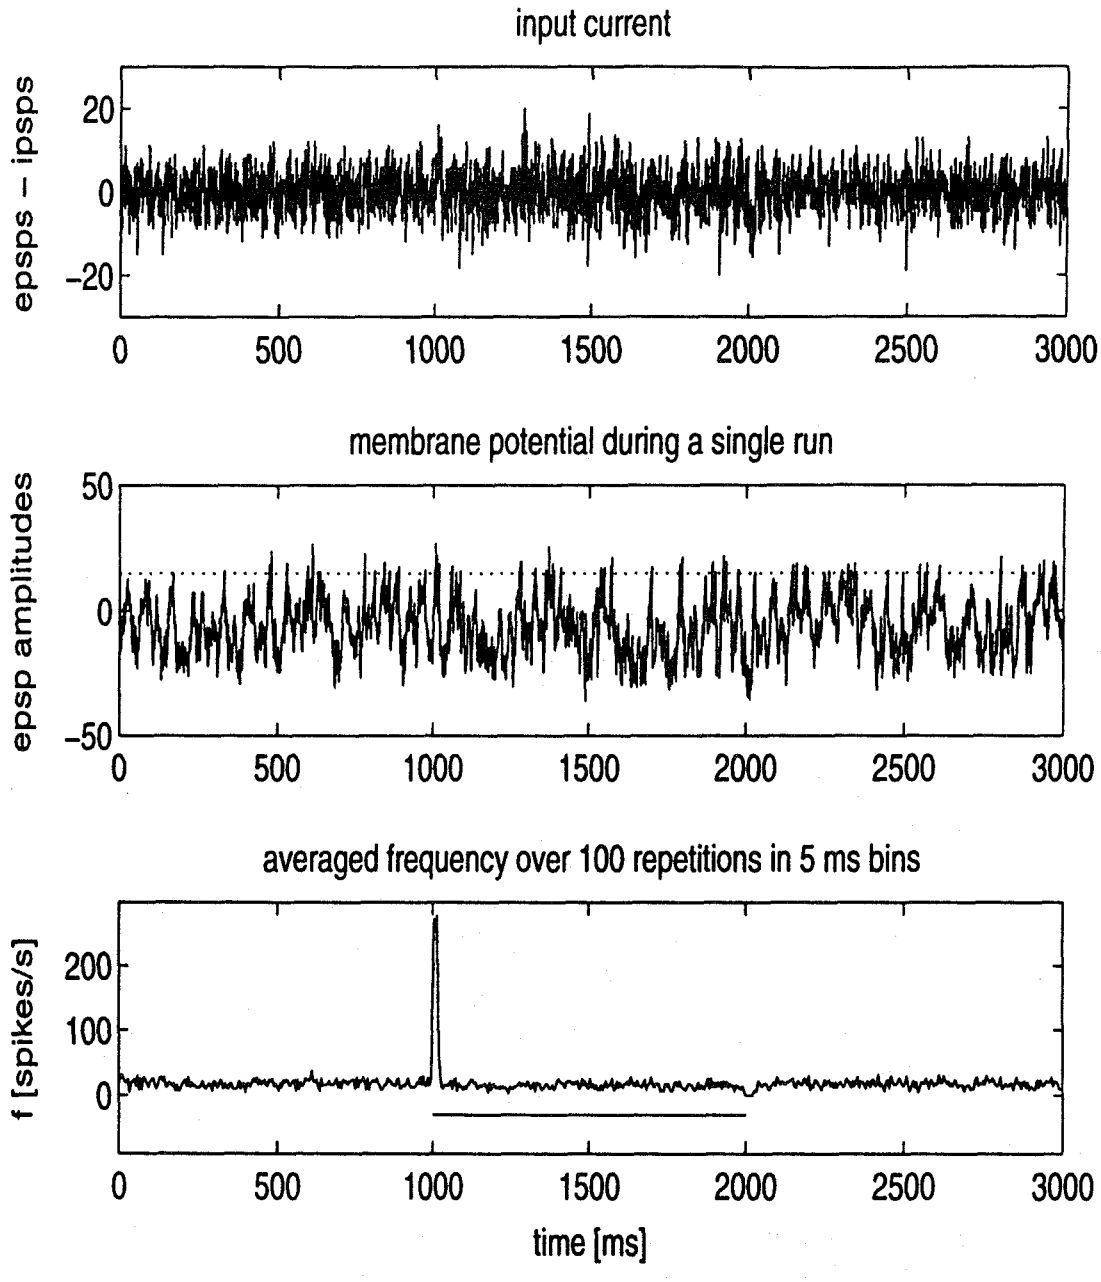

Figure 2: Single neuron model. (Top) The total input to the neuron is a sum of four Poisson processes, each with intensity $\lambda = 10$, consisting of excitatory and inhibitory background activity during the whole run and excitatory and delayed inhibitory signals ($\Delta^{\text{inh}} = 20$ ms) from $t = 1000$ to 2000 ms. (Middle) Membrane potential at the receiving neuron (time constant $\tau_{\text{m}} = 10$ ms); The dotted line indicates the firing threshold ($\theta = 15$). (Bottom) A spike histogram computed using 100 trials and 5 ms bins scaled to represent the mean firing rate in spikes per second. Note the pronounced onset response together with a rapid decay to the spontaneous rate even during stimulation. The stimulus duration is indicated by the horizontal bar.

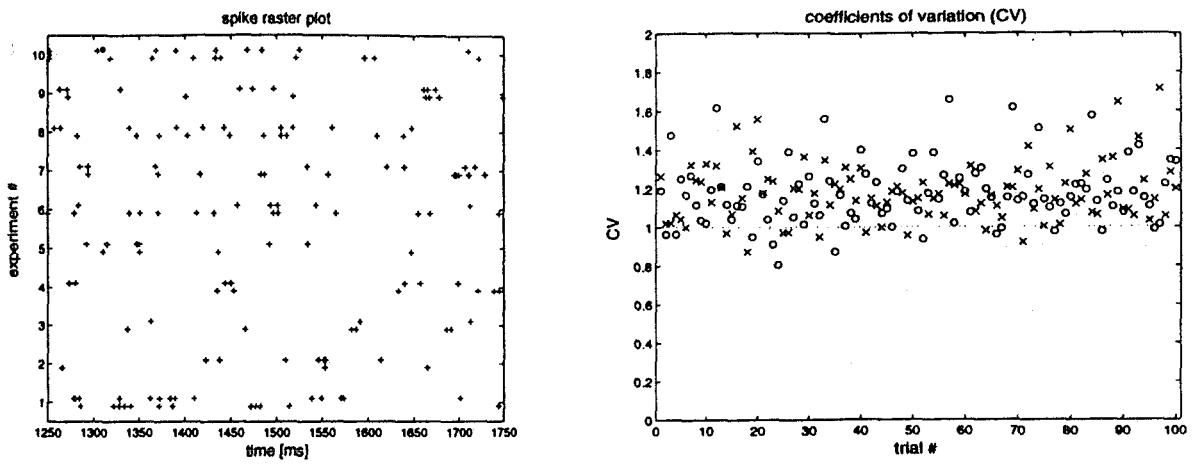

Figure 3: Spike trains for pairs of cells. (Left) Spike raster plot from 10 different runs of the same "experiment". Shown is the activity during stimulation of two cells respectively. Every plus denotes a single spike. The firing times are irregular as confirmed by the coefficient of variation shown to the right. (Right) Coefficient of variation (defined as the standard deviation over the mean of the inter-spike interval distribution) as a function of stimulus repetitions (the two cells are indicated by x and o). High values indicate high variability in firing. For a completely random spike train this value would be 1 (indicated by the dotted line).

where there was an excess of excitatory input due to the delayed arrival of the balancing inhibitory input. This can clearly be seen in the spike histogram (lower part of Fig. 2) obtained by averaging over 100 repetitions of the same experiment (but using a different seed for initializing the random number generator each time). The mean firing rate stayed constant throughout the whole run except for a pronounced burst at stimulus onset and a reduction of firing after stimulus offset. The large trial–to–trial variability in firing is shown in Fig. 3. On the left–hand side, single spikes from ten different trials are depicted while on the right–hand side a quantitative measure of variability, the coefficient of variation, is shown. Two cells are shown that share a common input, as described in the next section.

3.2 MULTIPLE CELL PROPERTIES

In the experiments of deCharms et al. (1995) simultaneous recordings of spike trains from pairs of cells were analyzed. We simulated two cells getting independent background signals but sharing identical stimulus components in their input ($C_1^{\text{stim}}(t) = C_2^{\text{stim}}(t)$ for all t). The top row of Fig. 4 shows correlations between the firing times of the two cells calculated for every single trial for three different periods of time (before, during, and after stimulus presentation) and averaged afterwards. There was a strong peak at zero time shift only during common stimulation indicating an increased tendency of the two cells to fire simultaneously.

This is not a surprising result, because one might expect the common input to drive both cells to firing threshold simultaneously, but it is worth noticing since only a fraction of the emitted spikes are affected. These synchronous spikes happen to occur at random times and are not stimulus–locked, as indicated by the flat shift predictor in the lower–right part of Fig. 4. The height of the central peak in the correlation depends mainly on the amount of common input relative to the total input to both cells as shown in Fig. 4 (bottom left).

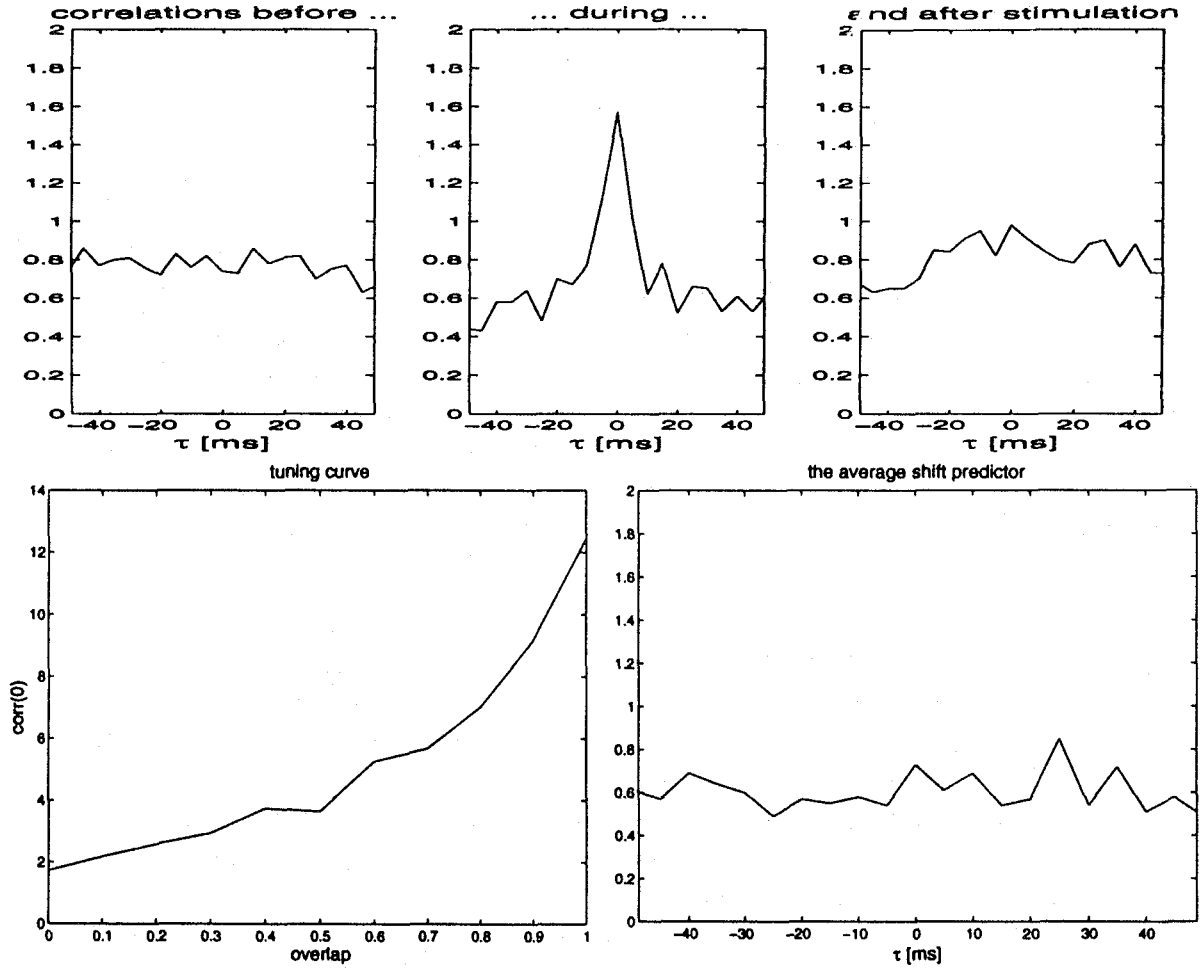

Figure 4: A pair of cells receiving a common input. (Top row) Average correlations from three different time periods: (Left) From $t = 250$ to 750 ms – before stimulation; (Middle) From $t = 1250$ to 1750 ms – during stimulation but after the onset response and (Right) from $t = 2250$ to 2750 ms – after stimulation. There is a clear peak at $\tau = 0$ for the stimulation period indicating that these two neurons have a tendency to fire in synchrony during presence of the stimulus. Correlations were calculated for every trial using 5 ms bins and averaged afterwards. (Bottom, left) Tuning curve: Height of the central peak in the correlation during stimulation as a function of the fraction of identical input. The peak height increases with the overlap. (Bottom, right) Due to the overall noisy structure of the observed response, the shift predictor, correlating responses from different trials, is flat. Thus, there was no stimulus–locked activity during the tonic phase of the response.

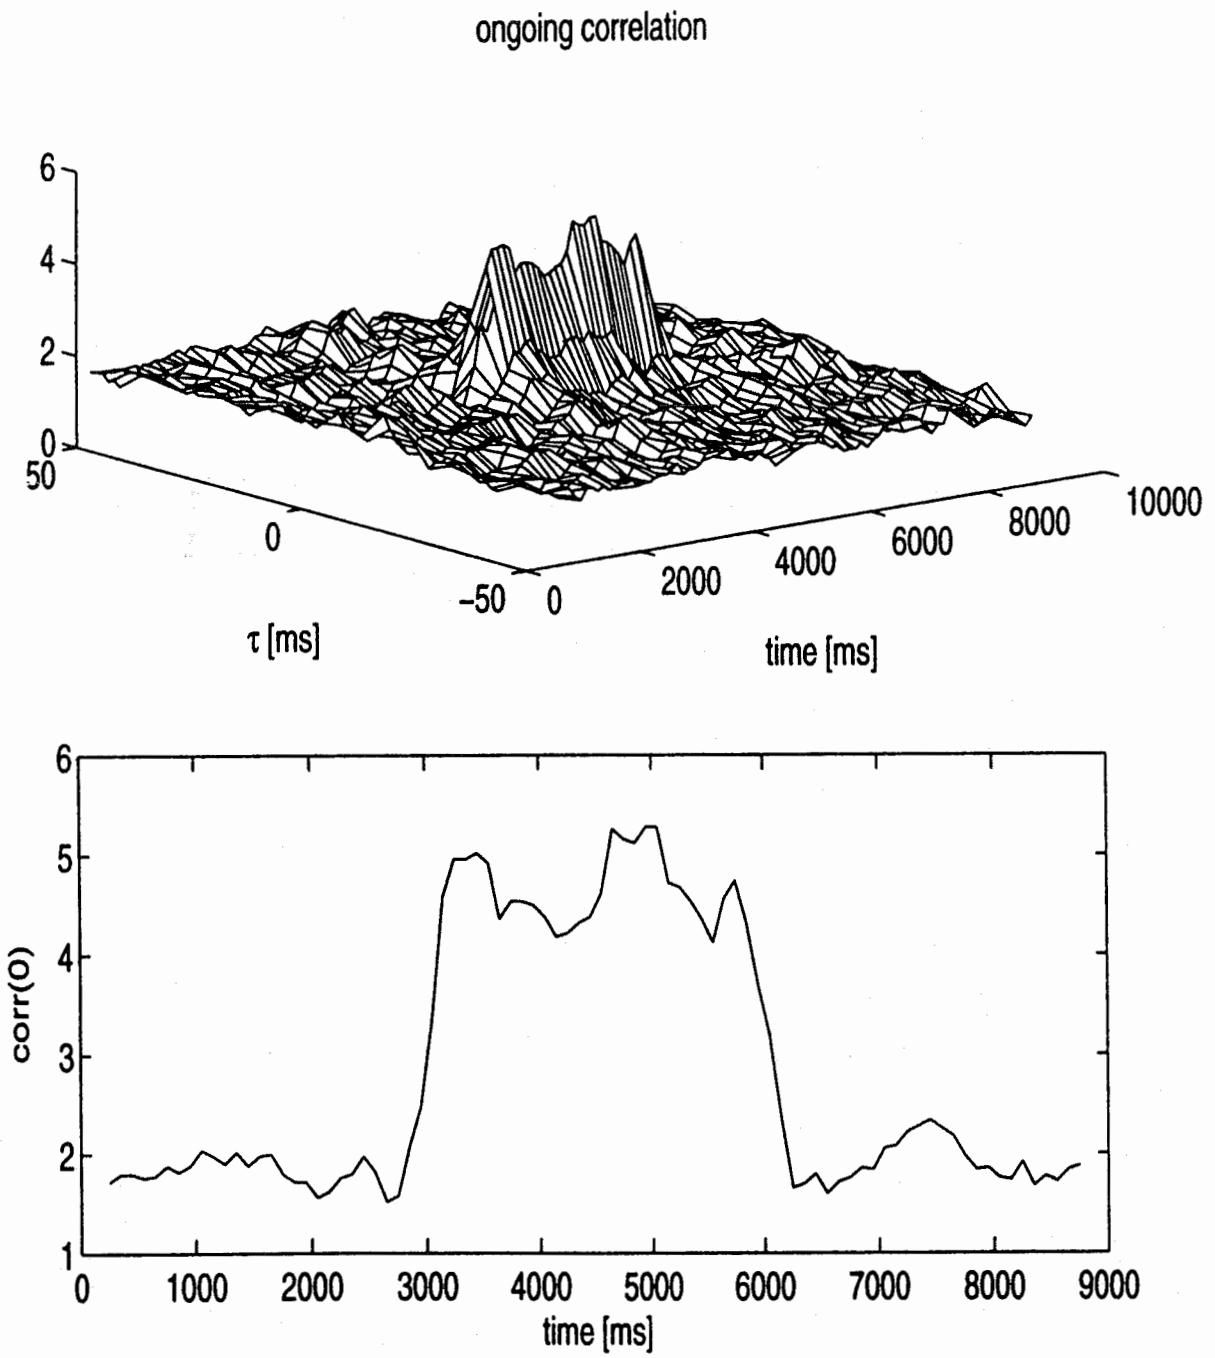

Figure 5: Time course of the average correlation calculated using a 500 ms time window sliding over a 9 s simulation run in 100 ms time steps. (Top) The full correlations as a function of time. (Bottom) The height of the central peak of the correlations as a function of time. A stimulus was presented from $t = 3000$ to 6000 ms.

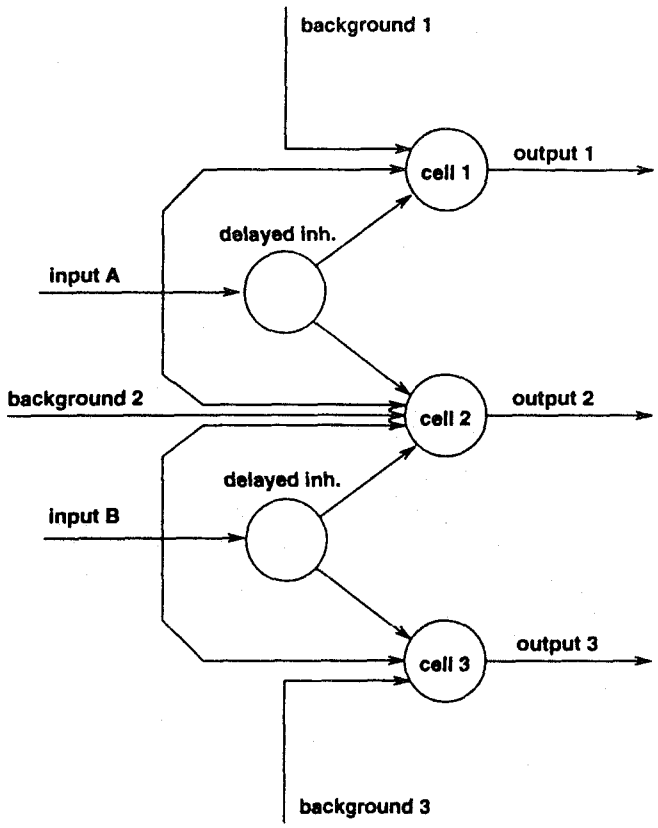

Figure 6: Basic structure for the three cell network. Now, cells 1 and 2 as well as 2 and 3 share some common input.

The overlap here is defined as the ratio of common versus total input, ranging from zero (no common input) to one (absolutely identical input).

Next, the time course of the correlation peak is shown in Fig. 5. Correlations were calculated from a 500 ms time window sliding over the entire run in 100 ms time steps leading to five times oversampling following deCharms et al. (1995). During the entire stimulation period, there was a pronounced increase in the correlations, which disappeared when the stimulus was turned off.

4 A FUNCTIONAL ROLE?

Finally, we suggest further computational implications of this mode of operation. It has been argued that the temporal structure of neuronal signals might be used for solving the binding problem (von der Malsburg 1981). Some time ago, there seemed to be experimetal evidence for this concept through the discovery of stimulus–related, collective oscillations, first found in the primary visual cortex of cats (Eckhorn et al. 1988, Gray and Singer 1989, Gray et al. 1989). Similar observations have also been made in monkeys (Kreiter and Singer 1996), but it is not clear whether the observed oscillations really have a crucial role in perception.

Here, we stress that the underlying mechanism for solving the binding problem is *simultaneaous* activity, not necessarily involving oscillations at all. Consider the neuronal network in Fig. 6. Two pairs of cells receive common input as before, so their output spikes show

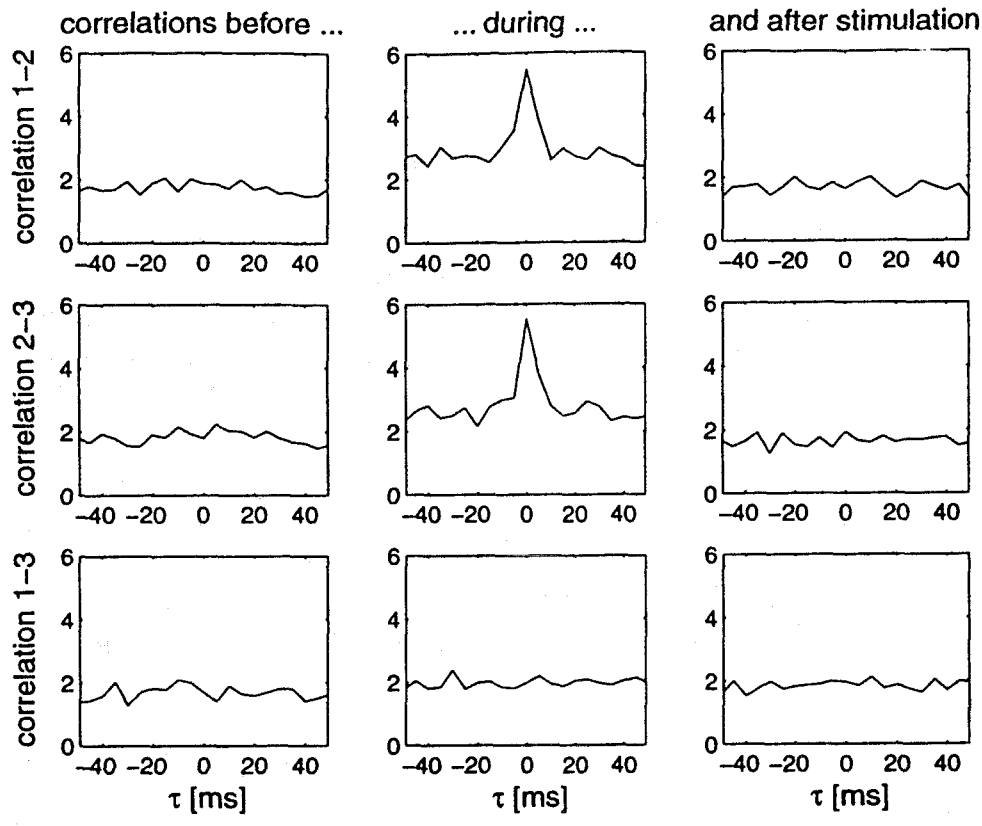

Figure 7: Demonstration of non–transitivity in a three neuron network. During stimulation, the responses of cells 1 and 2 as well as 2 and 3 show a tendency to fire in synchrony, but not cells 1 and 3. This is remarkable, since both are correlated with cell 2. This non–transitivity could be a useful property in avoiding the superposition catastrophe that can occur in binding the cell assemblies that represent multiple objects.

an increased tendency to appear simultaneously, as reflected in the correlations shown in Fig. 7. It is remarkable, however, that cells 1 and 3 are not correlated, despite the fact that these two cells both are correlated with cell 2.

5 DISCUSSION

In contrast to the common belief that neurons code information only in their mean firing rate, deCharms et al. (1995) have shown that there is another possibility of coding, based on the relative timing of spikes from different neurons. We have replicated their results in a neural model. Conceptually, this idea is not new, and the underlying firing pattern may be even more complicated than just synchronous firing, as in synfire chains (Abeles et al. 1993) or arbitrary firing patterns (Gerstner et al. 1993) or with respect to an internal neuronal clock (Hopfield 1995).

What is new here is the observation that relative spike timing might be used in a noisy mode of operation. For this regime, it has commonly been assumed that the only way to get at reliable information transmission should be based on a rate code (Shadlen and Newsome 1994). But there is increasing evidence for the possibility of temporal codes. First, it has been shown by Mainen and Sejnowski (1995) that neocortical neurons fire very reliably if

driven mainly by input fluctuations instead of a constant current. Therefore, the well known high variability in cortical spike firing times might reflect a high variability in the input to a neuron instead of intrinsic noise due to the spike generation process. Second, correlations in firing times between neurons tuned to similar stimulus features are omnipresent, but they have usually been interpreted as an artifact of common stimulation causing redundancy and having no use. Recently, this interpretation has been questioned. In the visual system, correlations seem to improve stimulus representation on the level of the retina (Meister et al. 1995) as well as the LGN (Dan et al. 1996). In the auditory system, deCharms et al. (1995) provided evidence for the crucial role of correlations in stimulus representations. Their study was the starting point for the model presented here. We do not claim to have reproduced every single detail of their data. For this, a biophysically more realistic model should be appropriate. But we have shown here how such a code might work naturally and reliably even in a noisy environment.

The final question, however, whether this type of coding is really used in the brain (i.e., read out at the next level) remains to be experimentally examined. Correlations are easily read out by neurons and they play a central role in learning, so there is every reason to continue along this line of investigation.

Acknowledgments

Supported by DFG (grant Ri 821/1-1) and The Howard Hughes Medical Institute.

References

Abeles, M., Bergman, H., Margalit, E., and Vaadia, E. (1993). Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *J. of Neurophysiology*, 70:1629–1638.

Dan, Y., Atick, J. J., and Reid, R. C. (1996). Efficent coding of natural scenes in the lateral geniculate nucleus: Experimental test of a computational theory. *J. Neurosci.*, 16(10):3351–3362.

deCharms, R. C., Schreiner, C. E., and Merzenich, M. M. (1995). Representation by spike coincidence in primary sensory cortex. *Soc. for Neuroscience Abstracts*, 21:1177.

Eckhorn, R., Bauer, R., Jordan, W., Brosch, M., Kruse, W., Munk, M., and Reitboeck, H. J. (1988). Coherent oscillations: A mechanism of feature linking in the visual cortex? *Biol. Cybern.*, 60:121–130.

Gerstein, G. L. and Mandelbrot, B. (1964). Random walk models for the spike activity of a single neuron. *Biophyhs. J.*, 4:41–68.

Gerstner, W., Ritz, R., and van Hemmen, J. L. (1993). Why spikes? Hebbian learning and retrieval of time–resolved excitation patterns. *Biol. Cybern.*, 69:503–515.

Gray, C. M., König, P., Engel, A. K., and Singer, W. (1989). Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature*, 338:334–337.

Gray, C. M. and Singer, W. (1989). Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proc. Natl. Acad. Sci. USA*, 86:1698–1702.

Hopfield, J. J. (1995). Pattern recognition computation using action potential timing for stimulus representation. *Nature*, 376:33–36.

Kreiter, A. K. and Singer, W. (1996). Stimulus–dependent synchronization of neural responses in the visual cortex of the awake macaque monkey. *J. Neurosci.*, 16(7):2381–2396.

Mainen, Z. F. and Sejnowski, T. J. (1995). Reliability of spike timing in neocortical neurons. *Science*, 268:1503–1506.

Meister, M., Lagnado, L., and Baylor, D. A. (1995). Concerted signaling by retinal ganglion cells. *Science*, 270:1207–1210.

Shadlen, M. N. and Newsome, W. T. (1994). Noise, neural codes and cortical organization. *Curr. Opin. Neurobiol.*, 4:569–579.

von der Malsburg, C. (1981). The correlation theory of brain function. Internal Report 81-2, MPI für Biophysikalische Chemie, Göttingen. Reprinted in: E. Domany, J.L. van Hemmen, and K. Schulten, Eds., *Models of Neural Networks II*, Chap. 2, Springer, Berlin, Heidelberg, New York, 1994