

Correlation Coding in Stochastic Neural Networks

Raphael Ritz and Terrence J. Sejnowski

Computational Neurobiology Laboratory
The Salk Institute for Biological Studies
10010 North Torrey Pines Road, La Jolla, CA 92037, USA

Abstract. Stimulus-dependent changes have been observed in the correlations between the 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. A simple neural model can reproduce most of the characteristics of these experimental observations based on model neurons having leaky integration and fire-and-reset spikes and with Poisson-distributed, balanced input. The source of the synchrony in the model was common sensory input. 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.

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 and Merzenich [3] 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:

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

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 $u_i(t) \geq \theta^{\text{inh}}$ 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, $C_i^{\text{bg}}(t)$ and $C_i^{\text{stim}}(t)$ respectively

$$C_i(t) = C_i^{\text{bg}}(t) + C_i^{\text{stim}}(t) , \quad (2)$$

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

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

where b denotes a balancing factor indicating the relative strength of the inhibition with respect to the excitation and Δ^{inh} represents a delay.

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

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

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 $\lambda = n_{\text{aff}} \cdot p_f$ 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 [4].

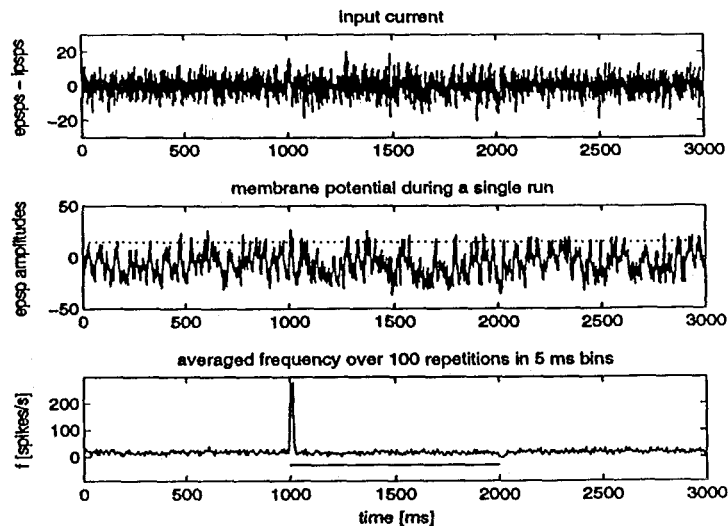


Fig. 1. Single neuron model. (Top) The total input to the neuron is a sum of four Poisson processes, each with intensity $\lambda = 10$, consisting of excitatory and inhibitory background activity during the whole run and excitatory and delayed inhibitory signals ($\Delta^{\text{inh}} = 20$ ms) from $t = 1000$ to 2000 ms. (Middle) Membrane potential at the receiving neuron (time constant $\tau_m = 10$ 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.

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 $\tau_m = 10$ ms and $\Delta^{\text{inh}} = 20$ 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. 1, the total input current fluctuates vigorously. The resulting membrane potential (Fig. 1, 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, 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. 1) obtained by averaging over




Fig. 2. 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.

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 where the experimental recordings show a transient increase of the firing rate. A large trial-to-trial variability in firing was observed as well (data not shown).

3.2 Multiple Cell Properties

In the experiments of [3] 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) assuming that cells with similar tuning properties are to some extent enervated by the same afferent neurons. The top row of Fig. 2 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




Fig. 3. 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. A stimulus was presented from $t = 3000$ to 6000 ms.

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. 2. 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. 2 (bottom left). The overlap here is defined as the ratio of common versus total input, ranging from zero (no common input) to one (absolutely identical input).

Finally, the time course of the correlation peak is shown in Fig. 3. 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 [3]. During the entire stimulation period, there was a pronounced increase in the correlations, which disappeared when the stimulus was turned off.

4 Discussion

In contrast to the common belief that neurons code information only in their mean firing rate, deCharms and Merzenich have shown that there is another possibility of coding, based on the relative timing of spikes from different neurons [3]. 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 [1] or arbitrary firing patterns [5] or with respect to an internal neuronal clock [6].

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 [9]. But there is increasing evidence for the possibility of temporal

codes. First, it has been shown that neocortical neurons fire very reliably if driven mainly by input fluctuations instead of a constant current [7]. 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 [8] as well as the LGN [2]. In the auditory system, [3] 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

1. M. Abeles, H. Bergman, E. Margalit, and E. Vaadia. Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *J. of Neurophysiology*, 70:1629–1638, 1993.
2. Y. Dan, J. J. Atick, and R. C. Reid. Efficient coding of natural scenes in the lateral geniculate nucleus: Experimental test of a computational theory. *J. Neurosci.*, 16(10):3351–3362, 1996.
3. R. C. deCharms and M. M. Merzenich. Primary cortical representation of sounds by the coordination of action-potential timing. *Nature*, 381:610–613, 1996.
4. G. L. Gerstein and B. Mandelbrot. Random walk models for the spike activity of a single neuron. *Biophysic. J.*, 4:41–68, 1964.
5. W. Gerstner, R. Ritz, and J. L. van Hemmen. Why spikes? Hebbian learning and retrieval of time-resolved excitation patterns. *Biol. Cybern.*, 69:503–515, 1993.
6. J. J. Hopfield. Pattern recognition computation using action potential timing for stimulus representation. *Nature*, 376:33–36, 1995.
7. Z. F. Mainen and T. J. Sejnowski. Reliability of spike timing in neocortical neurons. *Science*, 268:1503–1506, 1995.
8. M. Meister, L. Lagnado, and D. A. Baylor. Concerted signaling by retinal ganglion cells. *Science*, 270:1207–1210, 1995.
9. M. N. Shadlen and W. T. Newsome. Noise, neural codes and cortical organization. *Curr. Opin. Neurobiol.*, 4:569–579, 1994.