

Neurobiology Ionic Channels, Neurons, and the Brain

Edited by
Vincent Torre and
Franco Conti

NATO ASI Series

Series A: Life Sciences Vol. 289

BIOPHYSICAL ASPECTS OF CORTICAL NETWORKS

Stefan Rotter

Max-Planck-Institut für Entwicklungsbiologie
Spemannstraße 35/IV, 72076 Tübingen, Germany

1. INTRODUCTION

Artificial neuronal networks provide attractive models for cortical function, in particular, if “cognitive” properties emerge from their structure. Unfortunately, it turns out difficult to set up classical models which are comparable to the biological system on the level of single neurons. We look at artificial neuronal networks from a fresh perspective, which has the potential to extend their merits to a detailed and quantitative description of physiological phenomena in nerve nets of spiking neurons. In fact, the framework of stochastic point processes provides the tools for the construction of mathematically consistent models, which allow for a direct comparison with electrophysiological recordings on the level of individual nerve cells, in particular, if these are part of a large network. Moreover, the estimation of model parameters from experiments becomes feasible, so that a quantitative theoretical treatment as well as computer simulations of large networks under realistic conditions can be undertaken.

The enormous richness of dynamical behavior of cortical nerve nets is not least due to the fact that different mechanisms operate at different scales in space and time—cooperative phenomena involving a large number of nerve cells are tightly intermingled with subthreshold intracellular processes. To achieve a systematic account for the dynamical repertory of such networks with the help of a model it is of particular importance to find the right level of abstraction.

Action potential firing is the most visible result of signal processing in neurons like the pyramidal cells of the mammalian neocortex. The tendency of nerve cells to generate, under certain conditions, stereotyped rapid depolarizations across their membrane has led to the metaphor of an elementary decision making unit, which can only vote “all-or-none”. It is well established that different types of cortical nerve cells can generate temporal patterns of action potentials with very different characteristics. This diversity may come along with marked morphological or histochemical profiles. However, even more subtle variations in the dendritic geometry or only slight alterations in the ion channel equipment must be suspected to have impact on the dynamic repertory of a nerve cell (McCormick, Connors, Lighthall, Prince, 1985; Koch and Segev, 1989; Hille, 1992).

Action potentials are fundamental for the communication of a cortical neuron with other neurons in its spatial vicinity and in distant cortical areas, and with neurons in the sensory periphery or subcortical structures as well. In the cortex, the degree of convergence and divergence is on the order of 10^3 – 10^5 depending on the species and the cell type (Braitenberg and Schüz, 1992). Action potentials may travel over very long distances by the use of active transport mechanisms in the axon. Upon their arrival at a presynaptic terminal, transmitter substance is released into the synaptic cleft. This causes changes in the kinetic parameters of specific ion channels in the postsynaptic neuron, thus altering the cell's properties with respect to action potential firing. In this way, action potential firing in one cortical nerve cell can modulate the patterns generated by a great number of other cells. The details of signal transmission may differ from neuron to neuron, and their realization may also be correlated with well-established characteristics like the transmitter system used, or with less accessible differences in some biophysical or biochemical parameter.

In contrast to the highly reproducible waveform of individual action potentials, the behavior of a nerve cell has a strong stochastic component as far as spike timing is concerned. Most cortical neurons show a considerable variability in their response to repetitive presentations of identical inputs. This is well known from electrophysiological recordings in both single cell preparations and intact brains of behaving animals. At the extreme, spike trains from individual neurons in the frontal cortex, which are generated in distinct trials of trained behavior, may have no systematic structure at all, leading to a completely flat PSTH (Vaadia, Bergman, Abeles, 1989). In contrast, it has been demonstrated recently that the precision of action potential timing can be very much improved, if the neuron is driven by appropriate current transients (Mainen and Sejnowski, 1995). Little is known, however, about the relative contributions of the various sources of stochasticity in spiking activity. Even less is known about the brain's strategy to control the fluctuations, and about the functional consequences of their presence. The model we are going to employ, based on combined mathematical and biophysical reasoning, is hoped to contribute some understanding to these issues.

2. METHODS

The source of our reasoning is threefold: mathematical theory, electrophysiology, and numerical simulation. An intertwined sequence of arguments is taken from the theory of stochastic processes, in particular point processes, from neurophysiology and neuroanatomy, both empirical and theoretical, and from numerical simulations of stochastic models for large neuronal networks. Probabilistic arguments imply a certain general class of network models, which come along with specific experimental paradigms and a set of methods for statistical inference. The results of these physiological experiments are in turn fed back to theory and constitute the basis for further analytical arguments. Large scale computer simulations are performed by again using the parameters found in the experiments, and these finally connect back to the physiology of networks and brains. We give a short outline for the scope of each class of arguments.

2.1. Stochastic Point Processes

The theory of stochastic processes is a branch of mathematical probability theory, which is concerned with non-deterministic phenomena in general spaces, and with stochastic dynamical systems in particular. The notion of a Markov process is of paramount

importance, since it comprises a class of stochastic dynamical systems which can be successfully treated with analytical methods. Stochastic point processes render a probabilistic description for the distributed occurrence of point-like events—action potential firing in a biological neuronal network is a textbook example for a physical point process in space and time. The theory of such processes is very well developed and provides a concise framework for both the quantitative description and the statistical analysis of spiking activity. We refer to Daley and Vere-Jones (1988) for a contemporary introduction into the subject and for further guidance through the literature. An outline of our own contributions to a dynamic theory of interacting point processes and a point process theory of neuronal networks is deferred to the results section.

2.2. Slice Experiments in Rat Neocortex

The parameters of the stochastic model can be estimated from experiments *in vitro*, where one has tight control over a neuron's input and output without too much affecting its internal dynamics.

Experiments were made in acute coronal slices (400 μm) from visual cortex of 4–6 week old rats (Rotter, Heck, Aertsen, Vaadia, 1993). Using sharp glass micro-pipettes, we performed intracellular recordings from layer II/III pyramidal cells. The pipettes were filled with biocytin to allow for a morphological reconstruction of the recorded cells after the experiment. The experimental paradigm required that all sources of external inputs to the neuron be silent. To control for this, NMDA, non-NMDA, GABA_A, and GABA_B mediated ionic currents were blocked by adding the appropriate receptor channel blockers to the ACSF.

Substitute stimulation was presented to the neurons in terms of direct current input through the electrode. Constant depolarizing current at different strengths (up to 500 pA) was applied until the regular-spiking response of the neuron comprised a number of action potentials which was sufficient for a statistical analysis of the spike train (a few hundred). Several such spike trains, corresponding to different input currents, were the basis for the parameter estimation procedure as outlined in the results section. For more details of the preparation, the experimental paradigm, and the data analysis procedures we refer to Rotter (1994).

2.3. Numerical Simulations of Large Networks

A variant of the general stochastic point process model is adapted to cortical networks and implemented on a computer. Effective simulations can be performed in discrete time steps as well as in continuous time. Realizations of random variables like inter-spike intervals are obtained by the use of common pseudo-random number generators. Some technical details for the simulations are taken from the books by Ripley (1987) and Daley and Vere-Jones (1988).

The basic “physiological” parameters including amplitudes and time constants of after-hyperpolarization and postsynaptic currents are given random amplitudes and random initial values from a reasonably large range, as marked off by the experiments. No attempt is made to include specific topological structure into the simulated networks, the remaining “anatomical” parameters like synaptic convergence and divergence, and synaptic delays are also randomized in accordance with all available knowledge.

The simulations usually run over a few seconds or minutes physical time, where the variables recorded depend on the purpose of the simulation. In addition to a full record of

action potentials from a number of selected cells it is also possible to access subthreshold phenomena by recording “intracellularly” from one or more neurons. Each particular experiment is usually run several times, each time fixing different random values for selected parameters. Thus, it is possible to explore the dynamic repertory of such networks empirically and compare the simulation results directly to the outcome of electrophysiological experiments as, for instance, simultaneous extracellular recordings from several nerve cells in an intact brain.

3. RESULTS

3.1. Cortical Networks and Point Processes

Spatio-temporal patterns of action potentials. The function of an individual cortical neuron is given by its strategy of transforming many incoming spike trains into one single outgoing spike train. Only information encoded in streams of spikes can be faithfully transmitted to other neurons. From this perspective, signalling in terms of spatio-temporal patterns of action potentials is indeed fundamental to nervous coding in the cortex. For the time being, we distinguish the “objects” handled by such a network from the “operators” used by the network to manipulate its objects. The objects are spatio-temporal patterns of action potentials—they are the words used by the brain to represent and share information. The operators are given by the response properties of the compound network—their physical basis lies in the anatomy and physiology of its neurons, of the synaptic connections between them, and of the supporting structures. We will give more specific biophysical arguments and, at the same time, make this distinction mathematically explicit.

By all what is known from physiology, it is safe to idealize a cortical action potential as a point on the time axis, which has no temporal extension in itself, and which carries no amplitude information. Stochastic point processes then appear to be a suitable mathematical abstraction for a quantitative approach to spatio-temporal spike patterns in neuronal networks (Perkel, Gerstein, Moore, 1967). We claim that stochastic point processes are indeed fundamental for a theory of cortical function which accounts for the specific properties of a code relying on identical point-like events distributed in space and time, and which can in particular address the problem of controlling the degree of its stochasticity.

The occurrence of a single action potential is fully specified by the time epoch of its observation and the identity of the generating neuron. The combination of both data represents a point in space-time. A full realization of a point process is given by the complete record of all points generated during a specific experiment. It will be termed a “spatio-temporal pattern”. For a mathematical point process it is required that any finite time interval carries only a finite number of points. This, of course, is an obvious property of neuronal signals in physical networks. The process itself is fully specified by a probability law on the space of all admissible spatio-temporal patterns, with the additional requirement that counting points in any given time interval be a “measurable” procedure, in a mathematically well-defined sense. Thereby, the outcome of an observation of the process, which is a spatio-temporal pattern, turns into a random variable. Again, we refer to Daley and Vere-Jones (1988) for the mathematical details.

Temporal evolution of patterns. A point process in space and time can alternatively be considered as a stochastic dynamical system where point patterns gradually evolve in

time. This is achieved mathematically by successive conditioning on the internal history of the process, rather than looking at the probabilities for complete realizations. In fact, the full spatio-temporal pattern generated before a given time epoch can be viewed as the state of the system in that very moment. These states will be qualified as “external” to distinguish them from “internal” states, which will be introduced later. One obtains a Markov process, whose dynamics is given in terms of probabilities for transitions among states. The infinitesimal parameters of the process—then called point process intensities—are given by the probability for a point in one particular node normalized for an infinitesimal time interval (“response”), conditional on the pattern previously generated by all nodes (“input”). If these parameters exist, the resulting master equation is readily solved, see Appendix for more details.

The very nature of a Markov system is that its future evolution depends only on the present state, independent of the past. Incidentally, for a point process, each state includes a complete account of the internal history. The causal structure which is implicit in this description completely parallels the above statement that the function of the cortical network is defined by its strategy of transforming incoming spike trains into outgoing spike trains. In other words, the cortical hardware implements the set of transition probabilities, which have been termed operators above. From this point of view, the Markov approach to point processes is functional and therefore best suited for an interpretation in a biophysical context. This general framework sets the stage for a quantitative assessment of many otherwise untractable network phenomena.

Physical representation of states. Models for the function of biological neuronal networks are usually constructed bottom-up, on the basis of known biophysical properties of isolated nerve cells. They rely on the argument that the determinants of spike generation must also have some faithful representation in terms of physical properties of matter. The list of possibilities includes electrical properties of the cell membrane, metabolic states of its organelles, space-dependent concentrations of ions and neurotransmitters, conformational states of ion channels, and the like. It is hard to tell from the beginning which of the above should be considered for an adequate caricature of neuronal action and interaction. It is even less obvious which input features are actually represented in terms of physical states of single neurons. As a result, the potential of such models to grasp the essentials of neuronal network function can hardly be predicted and must be obtained from expensive numerical simulations.

The situation is different for the type of model we would like to put forward. The temporal evolution of any system of interacting point process has the spatio-temporal pattern of all previous points as a natural state variable. For each individual node, the point process intensity is a function of the external state, specifying the probability rate that it generates a point in response to the input. Thereby, the intensity functions provide a complete characterization for the system's stochastic dynamics, and any particular point process model amounts to making them explicit. The use of this concept for applications is only limited by the fact that the full space of patterns is somewhat unwieldy, it cannot be described by a finite number of parameters. For important examples of concrete point processes, including neuronal processes, it also seems to be unduly general and complicated as a physical system's state space, so that one may well ask for more specific, and simpler, models.

Mathematically, the spatio-temporal pattern of previously generated action potentials determines the instantaneous behavior of all neurons in the network. One cannot expect, however, that each individual neuron in a biological network with limited resources keeps

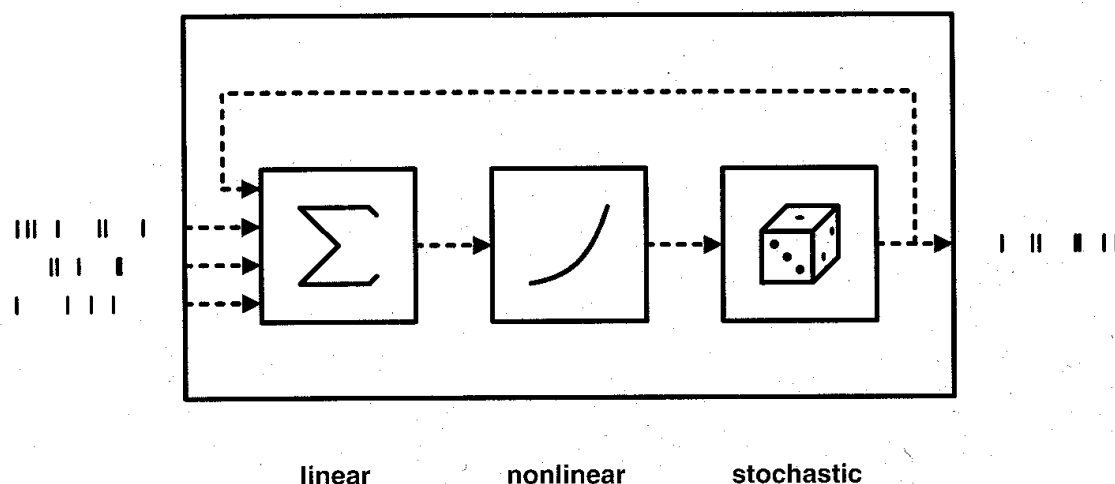


Figure 1. Schematic representation of single neuron function in the stochastic point process model. Inputs from other neurons in the network are first linearly integrated. The resultant internal state of the neuron is then non-linearly transformed into the point process intensity, which governs the stochastic generation of action potentials. The universal applicability of this scheme follows from pure mathematical reasoning without invoking physical or biological arguments. A biophysical interpretation of the parameters, however, leads to quantitative models of neuronal network function.

a full record of the previous history. For example, gross anatomy poses clear restrictions which forbid that the complete pattern is accessible to all neurons. On the other hand, some extent of information reduction is highly desirable for any device devoted to the processing of information. Elementary feature extraction, for instance, takes place in the processing of sensory information even on the level of single cortical cells, as is well known. In any case, the set of different physical states a single neuron can take at a given time instant may be expected to be small as compared to the set of all admissible external states. We will introduce additional features into the single neuron model to account for this.

"Integrate-and-fire" is universal. We discuss a formal way to achieve a model of reduced complexity by systematically accounting for redundancies in the space of external states, without appealing to biological mechanisms or relations among physical variables in the first place. In fact, purely algebraic manipulations on the space of patterns allow yet another interpretation of the probabilistic parameters of the system. For each neuron separately, we call two external states congruent, if they yield the same intensity in that neuron, as subpatterns of any larger pattern, and at all times. We collect all congruent external states into a single "internal state". Without going into the details of this construction, suffice it to say that one is led to a notion of internal states which can be guaranteed to be of reduced complexity, as compared to the external states. Nevertheless, internal states carry all information to fully determine the point process intensities of their nodes and, therefore, the stochastic dynamics of the system. The map which puts an internal state to its corresponding point process intensity will be termed "characteristic function". Finally, the projection from external states to internal states respects the "superpositionality" of patterns, which is important for the construction of models.

Each node now operates according to an "integrate-and-fire" scheme where the input pattern is first "linearly" projected onto some internal state and is then "nonlinearly"

transformed into the point process intensity. Figure 1 illustrates the idea. A small internal state space obviously narrows the dynamical repertory of the process. The initial problem is reduced to the question how the internal states are represented in terms of physical properties of the material nodes. Despite the fact that these states must remain unidentified in biophysical terms for the time being, one can safely pinpoint a number of correspondences between the biological network and its abstract point process model.

We assert that any internal single neuron state should correspond to some “minimal set of physical properties” of this particular neuron, which is sufficient to completely characterize the cell’s instantaneous readiness to generate an action potential. The projection of the external state onto the internal state corresponds to “dendritic integration”, where the personality of the neuron leads to a possibly very specific evaluation of the inputs obtained from other neurons. Each neuron’s dynamics also interferes with its own previous history. The superpositionality of internal states means, among other things, that any such state can be built from a less compound one by “adding” the effects of some extra action potentials. It may turn out that the temporal range of these “elementary” effects is finite or at least rapidly decreasing with time, as is typical for postsynaptic or after-hyperpolarization effects. In other cases, the induced changes in the neuronal behavior may be long-lasting or even permanent, an important example of this is long term potentiation and depression of synaptic gains upon appropriate stimuli. The resting state of the neuron may be identified with the “null-element” in the internal state space, because its effect on action potential firing equals the effect of the empty pattern, where definitely any input is absent. Finally, spike initiation, which is commonly thought to take place in the soma or at the axon hillock (Stuart and Sakmann, 1994), is governed by the point process intensity. The latter is obtained by evaluating the actual internal state using the characteristic function.

The class of models we have obtained as a mathematical consequence of very general assumptions represent only a slight generalization of some ad hoc models, which were invented to describe networks of synaptically coupled neurons and other biological systems (McCulloch and Pitts, 1943; Little, 1974; Hopfield, 1982; Mirollo and Strogatz, 1990). In these models, the input to a neuron is first evaluated by linearly adding all elementary contributions in terms of some internal state variable. Common views aim at the membrane potential or membrane currents, but it is not at all obvious whether these two are necessary or sufficient coordinates for a reasonable description of neuronal action potential firing. The neuron’s behavior is then determined by a non-linear transform of the internal variable, which specifies its instantaneous firing rate. Sometimes, this notion of rate coincides with our notion of intensity. In most cases, however, it has been interpreted as the expected number of spikes for a non-infinitesimal time interval. Then, no explicit action potential generation is assumed, and one clearly has no point process. In the classical case, all determinants of such an integrate-and-fire model must be constructed on the basis of indirect (and vague) external reasoning. In addition, common choices for the mapping of these operations on physical properties and mechanisms are biased by the methods of physiology.

3.2. A Stochastic Model for Cortical Pyramidal Cells

The situation for point process models is again different. Namely, in some cases, the physical nature of the abstract internal single neuron states can be identified with the help of experiments. As pointed out previously, any specific model amounts to an interpretation of the infinitesimal parameters in terms of biophysical properties of nerve cells. Each cell has its own state space with a linear structure and a nonlinearity transforming states into

rates. The rates—in contrast to the situation in some classical models—do not specify averages over some extended period of time. Rather, they are considered as instantaneous measures of neuronal excitation, very much like any other physical property of the neuron as the membrane potential or the number of open sodium channels. Braitenberg (1974) discusses an interpretation of conditional probabilities, and how they could be used to quantify neuronal action and interaction in the brain.

One is left with the task to estimate the time course of the point process intensity from experimental observations. Using the explicit knowledge of the solutions for the general master equation (see Appendix) one can obtain numerical fits for the parameters of integrate-and-fire models to experimental observations (Rotter, 1994). In fact, having available an estimate for the point process intensity as a function of time during an experiment opens the possibility to check the consistency of the implied hypotheses with the requirements of the model (see Figure 2).

By injecting direct current of different strengths and evaluating the systematic changes of the neuron's response, it is possible to decide on the physical nature of the internal state space of a regular-spiking neuron, and to give quantitative estimates of the parameters as well. Indeed, one can identify membrane current, rather than membrane potential, as the physical state variable containing more complete information on the neuron's state of excitation. Once this is established, it is possible to extract the explicit time course of (hypothetical) post-spike hyperpolarizing currents from the spike train statistics. Moreover, one obtains the approximate shape of the characteristic function translating the total membrane current into point process intensity, at least as far as it is covered by the data. Together with the substitute input current injected through the electrode, these data provide a complete stochastic characterization for a single neuron, according to the scheme represented by Figure 1. This simple parametric model is able to account for many details of spike generation upon direct current injection into regular-spiking pyramidal neurons, which constitute the major cell type in the neocortex (McCormick et al., 1985). The range of parameters covered by different subtypes of cells in different parts of the cortex, however, remains to be established. Finally, there is strong experimental evidence that this current-based description extends to synaptic interaction on the basis of spike-induced postsynaptic currents (Reyes and Fetz, 1993), which were pharmacologically blocked in our experiments. Generally, results of this type give new relevance to models of the integrate-and-fire type.

3.3. Toward a Network Model of the Cortex

Further analysis of the model equipped with parameters as extracted from experiments in real neurons should be expected to also yield information on the role of subthreshold processes for the emergence of collective behavior. For the present discussion, we pick three examples in support of this assertion.

Stability in cortical networks. A mathematical analysis of the “integrate-current-and-fire-spike” model and the corresponding stochastic point process contributes to the as yet unsettled problem of stability in recurrent neuronal nets. It turns out that a model network of regular-spiking nerve cells maintains a stable low level of activity, if all characteristic functions are non-decreasing, and if the matrix of synaptic couplings has a “dominant” diagonal. This is the case, for instance, whenever in each neuron the electric charge transported by after-hyperpolarization exceeds the total charge transported by excitatory and inhibitory synaptic currents. Both this assertion and the underlying hypothesis are in

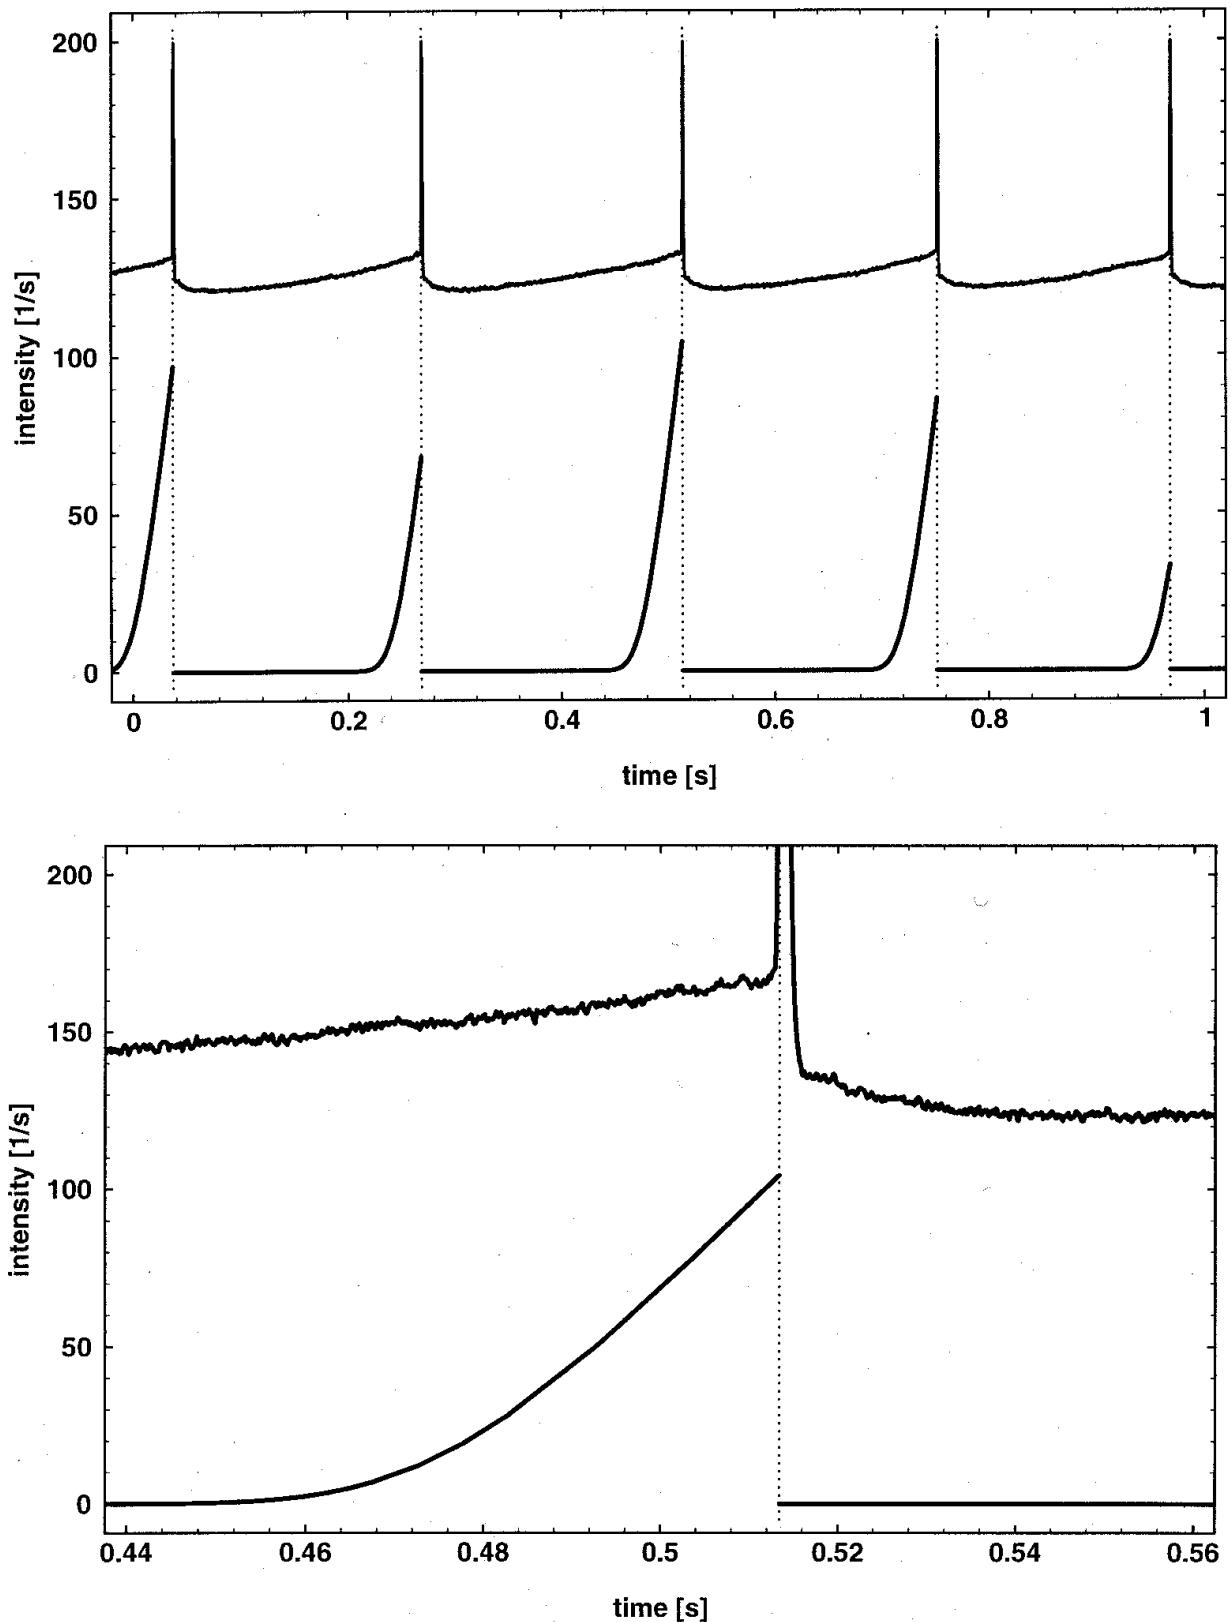


Figure 2. Empirical intensity of a neuronal point process from a current injection experiment. The intracellular recording (upper trace in both figures) shows the regular-spiking response to direct current stimulation, which is typical for cortical pyramidal cells. The corresponding point process intensity (lower trace in both figures) gives the probability for a spike normalized for an infinitesimal time interval. It is zero as long as the neuron is refractory or deeply hyperpolarized right after a spike, and it increases continuously until the action potential occurs. For the analysis, it is assumed that no spike but the most recent one influences the neuron's instantaneous excitability. The second figure shows the same data at a smaller time scale.

accordance with our knowledge of cortical physiology that synapses are sparse and post-synaptic currents are weak and fast, whereas after-hyperpolarization is strong and slow. Surprisingly though, inhibitory neurons are not required for stability, all activity control is taken care of by "auto-inhibition". Synaptic inhibition does contribute to achieve the low rates typically found in the cortex. An estimate of the maximal number of synapses which can be compensated for by after-hyperpolarization (10^2 – 10^5) is also in very good accordance with the anatomical figures (Braitenberg and Schüz, 1992).

Another interesting aspect of this parameter constellation is that the coefficient of variation for the inter-spike interval distribution of single neuron spike trains increases with the total amount of recurrent excitation in the network. Again, no inhibitory neurons are necessary to achieve coefficients as obtained from an analysis of *in vivo* data. This result is in contrast with some recent claims in the literature (Softky and Koch, 1993; Shadlen and Newsome, 1994).

Two time scales of synchronization. To organize and understand cooperative phenomena in larger networks, it turns out useful to characterize a single cortical pyramidal cell as a stochastic oscillator. The reason for this is not only its regular firing behavior upon direct current stimulation, but also its response to transient inputs. In fact, the susceptibility of the model neuron to synaptic inputs shows a distinct phase preference, leading to pronounced resonance phenomena. In addition, the characteristic function translating current into neuronal excitation turns out convex, which causes the neuron to prefer synchronous inputs over asynchronous ones. From this, one correctly predicts that transient synchronization of groups of neurons should be part of the dynamic repertory for recurrent networks of regular-spiking neurons.

Computer simulations of sparsely connected nets with random topology confirm that one can indeed distinguish two time scales of synchronization phenomena. At the macroscopic time scale extending over tens or hundreds of milliseconds, subgroups of neurons exhibit transient states of loose spike synchronization, with no appreciable effect on the average rates. Such "assemblies" desynchronize and reorganize themselves periodically upon persistent stimulation. Within the periods of enhanced group activity, one observes complex patterns of action-potentials extending in space and time, which recur with a precision in the millisecond range. This is all the more surprising as cell-cell interactions of the simulated networks were randomized in all their parameters including delay, amplitude, and time constant. Similar findings, both at the macroscopic time scale (Vaadia et al., 1995) and at the level of millisecond spike patterns (Abeles et al., 1993), have been reported from multi-neuron recordings in the prefrontal cortex of the behaving monkey.

On the plasticity of time structure. The ability of the cortex to keep its representations plastic are most valuable for the adaptive control of an agent's behavior in a highly variable environment. We began to investigate the possibilities of plasticity in time structure of neuronal signals by introducing physiologically inspired Hebb-like synapses into our stochastic networks. In the context of point process models with physiological parameters, this leads to a learning rule of the covariance type (see Appendix). From numerical simulations we conclude that the apparent weakening of synapses connecting out-of-phase neurons can serve to maintain a stable total amount of excitation within such networks. In contrast to most contemporary artificial neuronal systems in use, no artificial normalization of synaptic strengths is necessary. By employing a measure for the distance between spatio-temporal patterns, which naturally emerges from point process theory, we can demonstrate that the mere presence of plastic synapses has already useful consequences. One observes

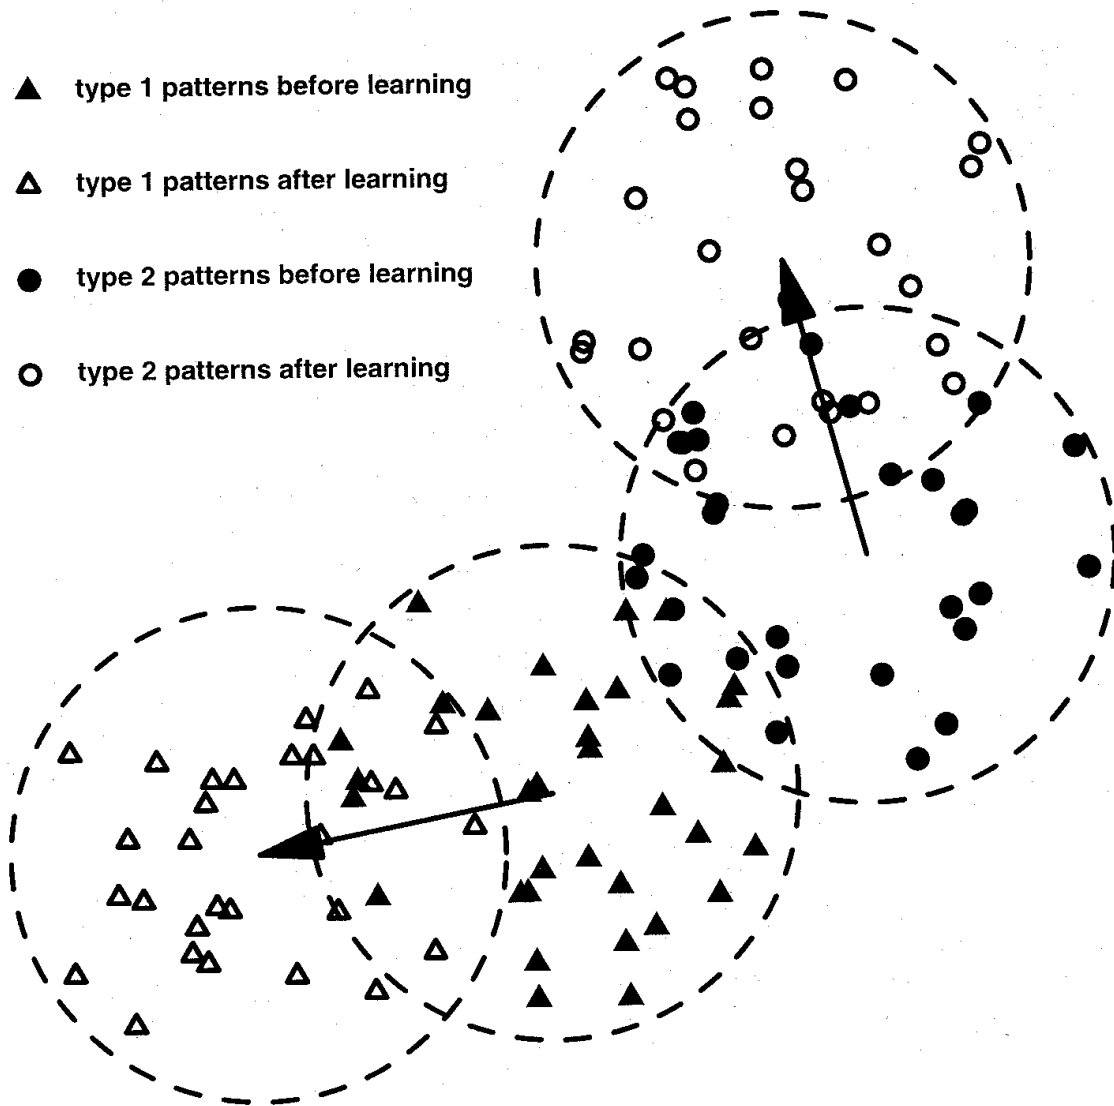


Figure 3. Representation of stimuli as spatio-temporal patterns. We show a schematic diagram which illustrates the effect of learning in a stochastic network. A small subnetwork responds to two different stimuli from surrounding neurons in two different ways (“type 1” and “type 2” responses) with a certain jitter and some overlap. This is indicated by the overlapping dashed circles comprising filled symbols. The mere presence of plastic synapses leads to a slow change of the representations for the two stimuli, which is finally non-overlapping while avoiding an increase in the jitter of the representation. This is indicated by the non-overlapping dashed circles comprising open symbols. The movement of the circles in space, as indicated by the arrows, stands for an increase in distance, or, equivalently, a decrease in similarity of the spatio-temporal patterns used for the representations.

the imprinting of spatio-temporal representations for temporally extended stimuli, as well as the automatic generation of discriminating spatio-temporal representations for several distinct stimuli (see Figure 3).

4. DISCUSSION

4.1. Summary

The spatio-temporal dynamics of cortical spike activity may be studied by using an abstract network model compatible with the notion of interacting stochastic point processes.

Any such system can be viewed as a Markov process, whose state at a given time instant is the spatio-temporal pattern of all previously generated spikes. The transition probabilities specify how a pattern gradually evolves in time. Neuronal network models which use a slightly generalized integrate-and-fire dynamics are a direct mathematical consequence of the assumption that neurons communicate by the use of action potentials. Assuming the existence of infinitesimal parameters, the corresponding dynamic equations are completely solved, and the solutions are used to identify important model parameters from electrophysiological recordings of real neurons. A simple parametric characterization of single neuron function is in fact achieved by fitting the model to the regular-spiking behavior of cortical pyramidal cells. A number of fundamental properties of recurrent cortex-like networks assembled from such neurons can be predicted, most notably their ability to maintain stable low rates of activity without the help of inhibitory neurons.

Computer simulations indicate that high precision spatio-temporal patterns, embedded in periods of enhanced cooperative group activity, may play a role for coding and computation in such networks. This is true, even if neither the anatomy of the network nor the physiology of its neurons are in any sense specially designed for that purpose. Plasticity of the temporal structure of such patterns is achieved by introducing Hebb-like synapses into the network. The resulting properties bring the point process model close to what more abstract neuronal networks are known to be capable of. By further exploiting our model system's known mathematical structure we expect to derive quantitative predictions which can also be applied to more complex experimental paradigms, involving, for instance, neural structures with specialized topology and plastic properties.

4.2. The Dynamic Repertory of Neuronal Processes

The notion of the "dynamic repertory" of a point process must remain vague until a mathematically well-defined formulation can be given. This is not possible so far, however, what we mean intuitively is something like the "true dimensionality of the system's phase space" or the "number of substantially different trajectories the process can take in a controlled way". Generally, the two determinants for the dynamic repertory of a system of interacting point processes are its dictionary, that is, the set of reachable spatio-temporal patterns, and its syntax, that is, the collection of rules governing the evolution of patterns. Clearly, both a more voluminous dictionary and a more flexible syntax increase the number of words at the system's disposal.

Among the many properties of the cortical network which contribute to its rich dynamic repertory we emphasize the sheer number of neurons, and the complexity of the network's organization as far as the processing of signals is concerned. The latter is reflected, among other things, by the degree of synaptic divergence and convergence, and by the degree of plasticity of the anatomical substrate during development. Another obvious source of dynamic diversity are cell-cell interactions, where extra complexity is added by postsynaptic currents with different time characteristics and possibly different reversal potentials, as well as by spike-induced modulatory currents. Yet another aspect, which is really at the heart of the stochastic model, is the precision in the control of timing as determined by the relationship between the strength of inputs and the stochastic properties of their integration. It is not unlikely that certain aspects of dendritic geometry and of the distribution of ion channels on the cell's soma and dendrite come into play at this point (Clay and DeFelice, 1983; Bernander, Douglas, Martin, Koch, 1991). Finally, we mention the non-stationary aspects of interactions among cells, the local or non-local characteristics of synaptic plasticity, and the depth and the time constants of memory. No doubt, all these

parameters will greatly affect the capacities of a neuronal network as an abstract language analyzer and generator.

In principle, all factors which contribute to the dynamic repertory of the neuronal process play a potential role for the evolution of cortical structure and function. One general evolutionary strategy could be to increase the number of distinguishable words in the dictionary to choose from, simply to provide for more powerful representations. This can be achieved by an increased number of independently operating mechanisms to elicit action potentials in single cells and mechanisms to synchronize spikes into spatial patterns which involve many cells. A related strategy could be to increase the flexibility and the adaptedness of the syntax to represent relations among objects. This might involve, for instance, plasticity of the interactions between cells by means of synaptic learning rules which evaluate correlations and similarities among the inputs.

5. APPENDIX

5.1. Point Process Dynamics

A spatio-temporal pattern is a list of points from space-time $\{1, \dots, n\} \times \mathbb{R}$

$$x = \{(i_1, t_1), (i_2, t_2), \dots\}.$$

such that any finite time interval carries only a finite number of points. A pattern at some time t has no points with time coordinates later than t . We consider the case where the probability for a transition of a pattern x at time s into another pattern y at time t is given by a density function

$$p_{s,t}(x, y)$$

which is then usually called point process likelihood. The point process intensity

$$\lambda_i^i(x)$$

represents the probability for a point at (i, t) normalized for infinitesimally small time intervals ("node i generates a spike in a small time interval around t "). The master equation of the process of patterns

$$\frac{\partial}{\partial t} p_{s,t}(x, y) = p_{s,t}(x, y) \left(- \sum_{i=1}^n \lambda_i^i(y) \right)$$

has the general solution

$$p_{s,t}(x, y) = \prod_{(u,j) \in y \setminus x} \lambda_u^j(y_u) \prod_{i=1}^n \exp \left(- \int_s^t \lambda_v^i(y_v) dv \right)$$

where y_u means the restriction of pattern y to times earlier than u . The explicit knowledge of these functions can be used for further theoretical analysis as well as for parameter estimation from experimental data.

5.2. A Synaptic Learning Rule

The following rule is used in simulations of stochastic networks

$$\Delta W_{ij}(t) = S_j(t)A_{ij} [W_{ij}^*(U_i(t)) - W_{ij}(t)]$$

where the parameters have the following meaning:

i	postsynaptic neuron
j	presynaptic neuron
W_{ij}	synaptic strength
ΔW_{ij}	change in synaptic strength
S_j	assumes 1 upon the arrival of a presynaptic spike, otherwise 0
A_{ij}	amplitude for the change in synaptic strength
W_{ij}^*	pivot for synaptic strength
U_i	postsynaptic depolarization

Depending on the depolarization of the postsynaptic neuron, this rule allows for both increase and decrease of synaptic efficacies, very much as found in hippocampal and cortical nerve cells (Artola, Bröcher, Singer, 1990). In the case of a non-decreasing transfer function it amounts to an iterative on-line measurement of the covariance between presynaptic and postsynaptic activity.

REFERENCES

1. Abeles M, Bergman H, Margalit E, Vaadia E (1993) Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *J Neurophysiol* 70:1629–1638
2. Abeles M, Prut Y, Bergman H, Vaadia E, Aertsen A (1993) Integration, synchronicity and periodicity. In: Aertsen A (ed.) *Brain Theory: Spatio-Temporal Aspects of Brain Function*. 149–181. Amsterdam, New York, London, Tokyo: Elsevier
3. Artola A, Bröcher S, Singer W (1990) Different voltage-dependent thresholds for inducing long-term depression and long-term potentiation in slices of rat visual cortex. *Nature* 347:69–72
4. Bernander Ö, Douglas RJ, Martin KAC, Koch C (1991) Synaptic background activity influences spatiotemporal integration in single pyramidal cells. *Proc Natl Acad Sci USA* 88:11569–11573
5. Braitenberg V (1974) On the representation of objects and their relations in the brain. In: Conrad M, Güttinger W, Dallin M (eds.) *Physics and Mathematics of the Nervous System*. 290–298. Berlin, Heidelberg, New York: Springer.
6. Braitenberg V, Schüz A (1991) *Anatomy of the Cortex—Statistics and Geometry*. Studies of Brain Function, vol. 18. Berlin, Heidelberg, New York: Springer
7. Clay JR, DeFelice LJ (1983) Relationship between membrane excitability and single channel open-close kinetics. *Biophys J* 42:151–157
8. Daley DJ, Vere-Jones D (1988) *An Introduction to the Theory of Point Processes*. Springer Series in Statistics. New York: Springer
9. Hille B (1992) *Ionic Channels of Excitable Membranes* (2nd ed.) Sunderland (MA): Sinauer
10. Hopfield JJ (1982) Neural networks and physical systems with emergent collective computational abilities. *Proc Natl Acad Sci USA* 79:2554–2558
11. Koch C, Segev I (1989) *Methods in Neuronal Modeling—From Synapses to Networks*. Cambridge: MIT-Press
12. Little WA (1974) The existence of persistent states in the brain. *Math Biosci* 19:101–120
13. Mainen ZF, Sejnowski TJ (1995) Reliability of spike timing in neocortical neurons. *Science* 268:1503–1506
14. McCormick DA, Connors BW, Lighthall JW, Prince DA (1985) Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex. *J Neurophysiol* 54:782–806

15. McCulloch WS, Pitts W (1943) A logical calculus of the ideas immanent in nervous activity. *Bull Math Biophys* 5:115–133
16. Mirollo RE, Strogatz SH (1990) Synchronization of pulse-coupled biological oscillators. *SIAM J Appl Math* 50:1645–1662
17. Perkel DH, Gerstein GL, Moore GP (1967) Neuronal spike trains and stochastic point processes. I. The single spike train. *Biophys J* 7:391–418
18. Perkel DH, Gerstein GL, Moore GP (1967) Neuronal spike trains and stochastic point processes. II. Simultaneous spike trains. *Biophys J* 7:419–440
19. Reyes AD, Fetz EE (1993) Effects of transient depolarizing potentials on the firing rate of cat neocortical neurons. *J Neurophysiol* 69:1673–1683
20. Reyes AD, Fetz EE (1993) Two modes of interspike interval shortening by brief transient depolarizations in cat neocortical neurons. *J Neurophysiol* 69:1661–1672
21. Ripley BD (1987) *Stochastic Simulation*. Wiley Series in Probability and Mathematical Statistics. New York: Wiley
22. Rotter S, Heck D, Aertsen A, Vaadia E (1993) A stochastic model for networks of spiking cortical neurons: Time-dependent description on the basis of membrane currents. In: Elsner N, Heisenberg M (eds.) *Gene—Brain—Behavior*. No. 491. Stuttgart: Thieme
23. Rotter S (1994) Wechselwirkende stochastische Punktprozesse als Modell für neuronale Aktivität im Neocortex der Säugetiere. *Reihe Physik*, vol. 21. Frankfurt: Harri Deutsch
24. Shadlen MN, Newsome WT (1994) Noise, neural codes and cortical organization. *Curr Opin Neurobiol* 4:569–579
25. Softky WR, Koch C (1993) The highly irregular firing of cortical cells is inconsistent with temporal integration of random epsps. *J Neurosci* 13:334–350
26. Stuart GJ, Sakmann B (1994) Active propagation of somatic action potentials into neocortical pyramidal cell dendrites. *Nature* 367:69–72
27. Vaadia E, Haalman I, Abeles M, Bergman H, Prut Y, Slovin H, Aertsen A (1995) Dynamics of neuronal interactions in the monkey cortex in relation to behavioral events. *Nature* 373:515–518
28. Vaadia E, Bergman H, Abeles M (1989) Neuronal activities related to higher brain functions—theoretical and experimental implications. *IEEE Transact Biomed Eng BME-36*:25–35

Acknowledgments

I would like to thank Dr. Detlef Heck (Washington University Medical School, St. Louis) for carrying out the slice experiments *in vitro*, which were the basis for parameter estimation in the stochastic model. I also thank Profs. Moshe Abeles and Eilon Vaadia (Hadassah Medical School, Hebrew University, Jerusalem) for providing the data from an *in vivo* experiment in the monkey VIGA, which entered the paper implicitly.