Dynamics of a recurrent network of spiking neurons before and following learning

Daniel J. Amit
Istituto di Fisica, Universita di Roma La Sapienza, Roma and Racah Institute of Physics, Hebrew University, Jerusalem

Nicolas Brunel
LPS, Ecole Normale Supérieure
24, rue Lhomond, 75231 Paris Cedex 05, France

October 23, 1996

Abstract

Extensive simulations of large recurrent networks of integrate-and-fire excitatory and inhibitory neurons in realistic cortical conditions (before and after Hebbian unsupervised learning of uncorrelated stimuli) exhibit a rich phenomenology of stochastic neural spike dynamics, and in particular, coexistence between two types of stable states: spontaneous activity, upon stimulation by an unlearned stimulus; and 'working memory' states strongly correlated with learned stimuli. Firing rates have very wide distributions, due to the variability in the connectivity from neuron to neuron. ISI histograms are exponential, except for small intervals. Thus the spike emission processes are well approximated by a Poisson process. The variability of the spike emission process is effectively controlled by the magnitude of the post-spike reset potential relative to the mean depolarization of the cell. Cross-correlations (CC) exhibit a central peak near zero delay, flanked by damped oscillations. The magnitude of the central peak in the CCs depends both on the probability that a spike emitted by a neuron affects another randomly chosen neuron and on firing rates. It increases when average rates decrease. Individual CCs depend very weakly on the synaptic interactions between the pairs of neurons. The dependence of individual CCs on the rates of the pair of neurons is in agreement with experimental data. The distribution of firing rates among neurons is in very good agreement with a simple theory, indicating that correlations between spike emission processes in the network are effectively small.

Key words: neural network model — integrate-and-fire neuron — spontaneous activity — spike rate distribution — learning — working memory — cross-correlations

Modeling the dynamical activity in cortex is becoming increasingly realistic as to allow direct confrontation with data of single and of multiple electrode recordings. The confrontation includes, beyond mean spike rates in spontaneous and in selective activity, also single cell spike statistics and cross correlations (Amit et al 1994, Softky and Koch 1993, Usher et al 1994, Tsodyks and Sejnowski 1995, Hansel and Sompolinsky 1996, Ben Yishai et al 1995). The question that naturally poses itself is: how much complexity of the neural elements and how much specific structure of the neural assembly are actually indispensable for the reproduction of the experimental data? Most of the work cited above has dealt with visual cortex, where data is more accessible and hypotheses about geometric connectivity structure are plausible. Moreover, in visual cortex neural dynamics is investigated in the presence of visual stimuli, which in turn relate to the assembly structure via receptive fields and tuning properties of cells.

We have been motivated by delay activity in associative cortex (Miyashita 1993, Wilson et al 1993, Goldman-Rakic et al 1990, Fuster 1995). Such delay activity could be sustained by the synaptic efficacies formed during training and hence its study relates directly to learning (Amit et al 1994, Amit 1995). In IT cortex, following an unlearned stimulus, the spike rates are all low, while a familiar stimulus (i.e. one used during training) generates enhanced delay activity in some 1% of the cells (Miyashita 1988). Both types of 'states' are maintained, in absence of the provoking stimulus, by the same set of synaptic efficacies in the same cortical module (column). In this area receptive fields are very large (Gross et al 1972). Yet, for delay activity, one does observe columnar compartmentalization (Miyashita 1993). The lack of obvious geometric organization allows raising questions about generic features of assembly dynamics. Apart from the intrinsic interest in the computational properties of associative cortex, it is of basic importance to know which features of neural dynamics, in vivo, actually depends on a higher level of modeling complexity.

We have recently shown that a theory of neural activity in a cortical module can produce the expected spontaneous activity and 'working memory' states at the level of mean spike rates, as well as the coexistence of the two (Amit and Brunel 1996, to be referred to as AB96). The theory predicts average firing rates in different neuronal populations, and in different network 'states', as a function of the parameters describing the model neurons and the 'anatomy' of the network, but there are other measured 'observables' which are not captured by it. In particular: Firing rate distributions, which are usually wide in cortex; Interspike interval, which in vivo exhibit a high degree of irregularity; Cross-correlations between spike emission times of different neurons, which often exhibit a pronounced peak near zero delay. In this article, we investigate the following questions: does a simple network of excitatory and inhibitory integrate-and-fire neurons exhibit these properties; which parameters control these 'observables'; and is a 'microscopic' dynamic network of interacting spiking neurons, in presence of correlations, well approximated by the theory of (AB96).

Preparation: Overview of the cortical module

'Anatomy' of the module

Theory (AB96) has shown that the crucial parameters governing the stability of spontaneous activity are the level of nonselective *stochastic* excitatory afferent on the neurons of the module from outside the module and the relative strengths of geometrically unstructured excitation and inhibition within the module. Both parameters are related to the columnar organization observed in experiment and underlying the model. Functional columns of the dimension of 1mm^2 parallel to the cortical surface appear in different parts of cortex. This corresponds to the anatomical findings that collateral connections between pyramidal cells as well as the ramification of the processes of inhibitory inter-neurons have this length scale in common (Braitenberg and Schütz 1991). In such a column there are about 10^5 pyramidal cells and some $2{\times}10^4$ interneurons.

Even when a column appears to be acting relatively autonomously, it remains an *open system* in two senses: First, it receives selective external afferents, either during the presentation of stimuli or from selective activity in other modules. Second, since at least one half of the afferent excitatory synapses of a typical pyramidal cell come from outside the local neighborhood (Braitenberg and Schütz 1991, Liley and Wright 1995), spontaneous activity in remote parts of the cortex alone, is a source of a high influx of spikes into a module (Bernander et al 1991, Rapp et al 1992, Amit and Tsodyks 1992, Shadlen and Newsome 1994, AB96). This second component we model by a random nonselective afferent to each synapse, representing the activity of neurons outside the module: a Poisson process with a rate equal to the mean spontaneous activity of excitatory neurons *inside* the module.

It represents a source of stochasticity in the dynamics of the module itself. The reason we model this afferent to be stochastic is an expression of our ignorance, not only of what the rest of the brain does, but also of the interactions of the brain with the outside world. Clearly, had all these been known, the system might have been considered deterministic, or almost so. At a still lower level there would be the stochastic effects due to temperature, which are an expression of yet another level of openness and unknown interactions with the outside world. This non selective flux is purely excitatory, because inhibitory neurons have local axonal ramifications, hence it has a fluctuating depolarizing effect on *all* cells in the module.

The other important parameter is local: the relative magnitude of overall *local* excitatory and inhibitory inputs to a cell. Stability requires (AB96) that local inhibition at least balance local excitation. A similar constraint appears in other studies (Shadlen and Newsome 1994, Tsodyks and Sejnowski 1995). Balancing excitatory and inhibitory inputs on the local level, in face of an inhibitory population deficit of a factor of 4 or 5, implies that the average inhibitory afferent must be significantly stronger than the excitatory one, as already noted in (Shadlen and Newsome 1994). This can be achieved in a variety of ways, which have all some empirical support: by more efficient synapses; by more reliable synapses or by higher inhibitory spike rates (Mountcastle et al 1969, McCormick et al 1985, Komatsu et al 1988, Lacaille

1991, Thomson et al 1993).

Network sizes and scaling

The simulations reported here were run on an IBM Risc work-station. This limits the size of the networks to 15,000 cells. Since afferent spike statistics and fluctuations are essential for reproducing activity distributions as seen in cortex, various quantities must be scaled to make up for the fact that the simulated networks are smaller than the estimated 100,000 cells in a cortical column. This we do in two ways:

- By increasing the probability of connection between neurons in the module;
- By increasing neural spike rates.

For example, in a network of 100,000 neurons, with connectivity level of 5% and average rate of $5\mathrm{s}^{-1}$, a neuron receives, on average, 250 spikes in 10ms from recurrent collaterals. Simulating a network of 10,000 neurons, one could use a connectivity level of 20% and an average rate of $12.5\mathrm{s}^{-1}$to obtain the same afferent spike number statistics. To check the effect on the network's dynamics of decreasing the connectivity level and the average firing rates, we have performed a small number of simulations with larger networks, some as large as 10^5 neurons.

Integrate-and-fire (IF) neurons

A simple model of a spiking neuron is the integrate-and-fire (IF) neuron (Knight 1972). It allows large scale simulations (thousands and even tens of thousands of neurons) and permits analytical estimates (see e.g. Tuckwell 1988, Amit and Tsodyks 1991, 1992, AB96). Such a neuron, though very simple, when receiving noisy inputs has temporal dynamics similar to those of a complex reconstructed pyramidal neuron (Softky and Koch 1993). Amit and Tsodyks (1992) have shown how to reduce a neuron with arbitrary geometry, governed by non-linear cable equations, to an IF neuron with renormalized parameters. Thus it may not be necessary to model the individual neuron in great detail for the restricted set of observables at hand. In many cases the simplified models and more detailed ones reach similar conclusions regarding network dynamics. To what extent identifiable dynamical behavior of single neurons in a cortical environment depends on detailed ion-channel dynamics is one of the subjects we shall try to address in this study, through a comparison of our results with more complex models.

Synaptic structure – sources of disorder

Neurons are connected by synapses, of which the local collateral ones are explicitly modeled. Whether the network is in its '*vacuous*' state, in which it can sustain only nonselective spontaneous activity, or in its *instructed* state, in which synapses have been modified by learning from stimuli, and structured delay activity can be sustained, there are several sources of disorder in the synaptic structure. They play an

essential role in what follows. The main division line of the synaptic sources of disorder is between disorder which is fixed (*quenched*) on time scales long compared with the internal neural time constants and fast-changing (*annealed*) for which synaptic efficacies can fluctuate on very short time scales. The model contains three types of disorder, which can be each quenched or annealed:
1. 'Transmission probability': the probability ϵ that a spike emitted by a neuron reaches another randomly chosen neuron. Its quenched version is the 'connectivity level' c, i.e. the probability of forming a (fixed) synapse between any pair of neurons in the column. The annealed version is 'reliability', i.e. the probability r that a pre-synaptic spike arriving at a given existing synapse generate a post synaptic potential (PSP) (See e.g. Thomson et al 1993a,b, Tsodyks and Sejnowski 1995). The transmission probability is a product of these two parameters, $\epsilon = cr$.
2. Random efficacy distribution: The efficacies vary from synapse to synapse.
3. Synaptic transmission delays: Each existing synapse has an associated delay between a pre-synaptic spike and the PSP produced by it. These delays vary from synapse to synapse with a given average and variability. Both synaptic efficacies and delays are either fixed (expressing the intrinsic variability of synapses due for example to their different positions on the dendrite) or fluctuating during the operation of the network (expressing fast synaptic noise).

Summary of results

Two modes of operation

The network was studied in two modes of operation:

1. Prior to learning: synapses are totally random and we expect spontaneous non-selective activity.

2. Following Hebbian learning of external stimuli: Excitatory synapses potentiated by stimuli on top of the previous random distribution. We expect unstructured spontaneous activity following stimulation by an unfamiliar stimulus; and stimulus-selective delay activity following stimulation correlated with one of the previously learned stimuli (AB96).

Unstructured network

1. The distribution of firing rates is rather wide, due to the variability in the number of connections from neuron to neuron. This distribution is in very good agreement with theoretical predictions.

2. The distribution of the interspike intervals (ISI) in the emission process of single cells is similar to a Poisson process, except at small intervals. The variability in the spike emission process is high for low rates and is effectively controlled by the value of the post-spike reset potential relative to the average depolarization of the membrane, as in (Tsodyks et al 1994).

3. Network dynamics depends very weakly on the type of 'implementation' of synaptic disorder, at parity of 'transmission probability': the behavior of a fully reliable synaptic matrix with random connection probability ϵ is very similar to a fully connected synaptic matrix with reliability ϵ;
4. The global activity in the network has very large, short lived, fluctuations. The total number of emitted spikes increases several times above average and is then suppressed by rising inhibition.
5. The spike trains of pairs of neurons exhibit a pronounced peak at zero delay in their CCs, related to the fluctuations of the global activity.
6. The size of the peak in the average CCs relative to the background increases when average firing rates decrease, at fixed transmission probability;
7. In the range of parameters studied, the size of the peak in the CC has small influence on mean spike rates. Hence, the existing correlations between spike processes on different cells are effectively small. We expect this to be also the case in a cortex;
8. CCs of individual pairs of neurons depend very weakly on synaptic efficacies connecting the pair: they are essentially indistinguishable for two neurons that are mutually connected, have a one way connection or no connection.
9. Individual CCs depend more strongly on firing rates of the two cells: neurons with lower rates are typically more correlated than neurons with high rates.
10. CC's are sensitive to the ratio of external to recurrent currents, because the external currents decorrelate the activity in the module.
11. Random spike transmission delays have little effect on the average rates in the network. They have some effect on the temporal structure of the global activity fluctuations, reflected in the average CC's: The time constant of the damped oscillations observed in the CCs increases with the average delay, while the amplitude of the central peak slightly decreases.

Networks structured by learning - selective activity

Detailed description of network dynamics following Hebbian learning will be presented in a future communication. Here we limit ourselves to an example. Learning consists of stochastic LTP in synapses among pairs of excitatory neurons simultaneously driven by a stimulus and LTD among neurons of which one is driven and one not.

1. If potentiation is not strong enough, no selective delay activity appears, but spontaneous activity becomes inhomogeneous: *all* neurons which responded to any one of the stimuli increase their rates, while other neurons decrease their spontaneous rates.

2. Beyond a critical potentiation level, the spontaneous activity state *coexists* with a multiplicity of selective 'working memory', or delay activity, states. The spontaneous activity distribution is excited by any stimulus different enough from all the stimuli previously learned (see e.g. Fig 1 of Miyashita 1988). Following a stimulus similar to one of the learned stimuli, the network maintains a stationary distribution of spike rates in which a small selective group of cells has significantly elevated rates which persist following the removal of the stimulus (delay activity). Most of the neurons in the module emit spikes at low rates.

3. The average firing rates of neurons which participate in a given excited state can be relatively close to the spontaneous rates (typically there is a 4:1 to 10:1 ratio). These ratios are similar, though slightly higher, than those observed in the delay experiments in primate cortex (see e.g. Wilson et al 1993, Nakamura and Kubota 1995).

These phenomena are well captured by theory (AB96).

Methods - Details of the model

Network architecture

We study a network composed of N neurons, from which $N_E = 0.8N$ are excitatory and $N_I = 0.2N$ inhibitory. Each neuron in the network has a probability c (connectivity level) of a direct synaptic contact with any other neuron (excitatory or inhibitory). Hence each neuron receives on average cN_E recurrent excitatory synapses, and cN_I recurrent inhibitory synapses. This leads to an amorphous recurrent synaptic structure. Each neuron receives an unselective stochastic afferent representing the synaptic input from spontaneous activity in remote parts of cortex. In the simulations N varied between 6,000 and 15,000 for networks with fixed (quenched) synapses and up to 100,000 for 'annealed' networks (see below) and 'transmission probabilities' ϵ between 0.05 and 0.25.

Neurons

Excitatory and inhibitory neurons are modeled as IF cells. Neuron number i of type α ($i = 1, \ldots N_\alpha$; $\alpha =$ E,I) has, at time t, two dynamical variables associated with it: a depolarization at the soma $V_i^\alpha(t)$ and an afferent post-synaptic current at the somatic integrator $I_i^\alpha(t)$. It is parameterized by an integration time constant τ_α; a threshold θ; a postspike reset potential V_r and an absolute refractory period τ_0. (The last three are taken equal for all neurons.)

The membrane depolarization of a neuron of type α, below threshold, is modeled by:

$$\tau_\alpha \dot{V}_i = -V_i + I_i \qquad (1)$$

with the synaptic current I_i expressed in units of the potential. When V_i reaches the threshold θ, the neuron emits a spike and its potential is reset to V_r, following an absolute refractory period τ_0. I_i is the sum of the recurrent synaptic input I_i^{rec} and the external synaptic input I_i^{ext}. Assuming that the synaptic conductance dynamics is fast relative to the neural integration time, the net *recurrent* depolarizing synaptic input charging the membrane of neuron i at time t is (see e.g. AB96):

$$I_i^{rec}(t) = \sum_{j \in E,k} \rho_j(t) J_{ij} \tau_\alpha \delta(t - \tau_{ij} - t_j^k) - \sum_{j \in I,k} \rho_j(t) J_{ij} \tau_\alpha \delta(t - \tau_{ij} - t_j^k) \qquad (2)$$

where the reliability variable $\rho_j(t)$=1 with probability r and is 0 otherwise, independently at every synapse and for every spike (see also Tsodyks and Sejnowski 1995). J_{ij} is the synaptic efficacy from neuron j to neuron i, τ_{ij} is the delay associated with this synapse, and t_j^k is the time neuron j emitted spike number k. The summations are over $j \in E, (I)$, i.e. over excitatory (inhibitory) presynaptic neurons and over the spikes k, arriving at a particular synapse. J_{ij} is the amplitude of the PSP provoked in neuron i by a spike arriving at this synapse. The synaptic variables will be defined in more detail below.

We take τ_E=10ms; τ_I=5ms, to express the fact that inhibitory neurons are typically faster than excitatory ones. θ=20mV above resting potential; V_r=10mV (except in Section 'Single cell spike statistics'); τ_0=2ms.

Recurrent synaptic interactions

Synapses in the network are of four possible types, as each of the presynaptic and postsynaptic neurons can be either excitatory or inhibitory. A synaptic quantity labeled $\alpha\beta$, connects a presynaptic neuron of type β to a postsynaptic neuron of type α. Synapses are drawn at random according to one of the following procedures:

A. **Fixed random connectivity**: Every neuron receives an equal number of contacts, from cN_E excitatory and cN_I presynaptic neurons, chosen at random.

B. **Variable random connectivity**: A synapse between any pair of neurons is present with probability c *independently* for all pairs of neurons. The number of contacts a neuron receives is a binomial random variable with mean cN_E (cN_I) and RMS $\sqrt{c(1-c)N_E}$ ($\sqrt{c(1-c)N_I}$), for excitatory (inhibitory) contacts, respectively.

In the simulations both procedures have been tested. Procedure A is best suited for comparison with the previous analysis (AB96), since it makes all neurons equivalent with respect to the connectivity scheme. The analysis can be extended to include variability in the connectivity (see Appendix B). The efficacy of the synapse transmitting from neuron j to i is denoted J_{ij}. If (ij) are of types $\alpha\beta$, its value is chosen from a Gaussian distribution ρ_J with mean $J_{\alpha\beta}$ and width $J_{\alpha\beta}\Delta$. The delay associated with synapse (ij), τ_{ij}, is chosen at random from a uniform distribution ρ_d with mean τ_d ranging between $(1-s_d)\tau_d$ and $(1+s_d)\tau_d$.

Finally, there may be fast (annealed) synaptic disorder, i.e. 1. Each synapse has reliability r: when presynaptic neuron j emits a spike, it provokes a PSP in the postsynaptic neuron i, of magnitude J_{ij} (0), with probability r (1–r), respectively. 2. When a spike arrives at a postsynaptic neuron, its efficacy and delay are drawn randomly from distributions ρ_J and ρ_d respectively, independently for different spikes. This type of disorder allows for the simulation of larger networks.

We have simulated networks in the following parameter ranges:

- Synaptic efficacies:
  - Excitatory to excitatory, J_{EE}, 0.1–0.5mV;
  - Inhibitory to excitatory, J_{EI}=3J_{EE} (0.3 to 1.5mV).
  - Excitatory to inhibitory, J_{IE}=2J_{EE} (0.2 to 1mV).
  - Inhibitory to inhibitory, J_{II}=3J_{IE}.
- Variability of synaptic efficacies: Δ, in the range 0 to 0.5.
- Spike transmission delay: τ_d, average 1ms, s_d =0.5.
- Synaptic reliability: r, in the range 0.05 to 1.

Two limit cases:

- Fixed (quenched) synaptic structure: Connectivity level c varies from network to network. For each network the synaptic system and the delays are fixed, i.e. reliability parameter $r = 1$.
- Fully 'annealed' synaptic structure: The reliability parameter r varies from network to network. The connectivity level $c = 1$ (all to all connections). Both synaptic efficacies and delays are selected at random and independently at each step. The entire synaptic structure fluctuates in time and is uncorrelated at different time steps.

Non-selective external input

Each neuron in the local module receives a non-selective, stochastic external input. It represents the background activity of other cortical areas that provide inputs to the network. The rates of the excitatory neurons in these areas are assumed to be similar, on average, to the rates of the neurons in the local network. Furthermore, the efficacies of the synaptic connections from outside the network are assumed to be, on average, equal to the recurrent excitatory connections on the same neurons.

If no stimulus is presented to the network, the total external current to neuron i between t and $t+dt$, $I_i^{ext} = I^\star$, is chosen randomly and *independently* at each neuron and each time step, according to

$$I_i^\star(t) = J_{\alpha E}\mathcal{N}_i(t)$$

where $\mathcal{N}_i$ is drawn from a Poisson distribution with mean $\epsilon N_E \nu_E dt$, and ν_E is the average spontaneous spike rate of excitatory neurons within the module. In other words, the non-selective external input to a cell is equal, on average, to the input it would receive, in the same time interval, from local collaterals that are firing at spontaneous rates.

Structured external input – external stimuli

For simplicity each of p stimuli is characterized by an N_E-bit word, η_i^μ=0,1 ($i = 1, \ldots, N_E$, $\mu = 1, \ldots, p$), in which the 1-bits define a subset of the excitatory neurons which are activated by stimulus μ. Each stimulus activates a random, and uncorrelated, subset of fN_E excitatory neurons. f is the coding level, or sparseness, of a stimulus in the module. When stimulus number μ is presented, the external currents I_i^{ext} become

$$I_i^{ext}(t) = I_i^\star(t)(1 + H\eta_i^\mu) \tag{3}$$

where H is the selective 'contrast' of the stimulus with respect to the background. The response of a cell during the presentation of a stimulus is its 'visual response'.

H is typically small, between 0.1 and 0.2. It is spontaneous activity, which maintains neurons most of the time depolarized near threshold, that produces a sensitivity to stimuli with low contrast (see also Tsodyks and Sejnowski 1995).

Learning

The presentation of stimuli induces modifications in the synapses of the unstructured network, whose behavior is simulated and described first. Upon learning, starting from its initial state J_{ij}, a synaptic efficacy may either remain unmodified, or make transitions to one of two possible states: a potentiated state $J_{ij}+\delta J_+$, and a depressed state $J_{ij} - \delta J_-$. Thus any synapse has three possible states; the state eventually chosen depends on the history of stimuli presentations. Possible transitions between these states are

- Potentiation (LTP), the transition from either J_{ij} or $J_{ij} - \delta J_- \rightarrow J_{ij} + \delta J_+$, occurs with probability q_+ if both pre- and postsynaptic neurons are activated by a given stimulus (i.e. $\eta_i = \eta_j = 1$ for the particular stimulus);
- Depression (LTD), the transition from either J_{ij} or $J_{ij} + \delta J_+ \rightarrow J_{ij} - \delta J_-$, occurs with probability q_- if only one out of pre- and postsynaptic neurons is activated by the stimulus.

Otherwise no transition occurs. (See e.g. Amit and Fusi 1994, Amit and Brunel 1995a). Note that once a synapse has made a transition it cannot go back to its initial state J_{ij}. Such simplified dynamics schematizes roughly experimental observations about synaptic plasticity in neocortex (see e.g. Bliss and Collingridge 1993, Artola and Singer 1993). The synaptic dynamics is not explicitly modeled. Instead, we consider a situation in which the p stimuli have already been presented a large number of times (overtraining) in a random sequence. In this case, the probability distribution of the resulting synaptic efficacies as a function of the p stimuli, can be explicitly computed (Brunel et al 1996, to be published). The result is given in Appendix C. In the resulting synaptic matrix, potentiated efficacies have an average value of $J_1 = J_{EE} + \delta J_+$, and depressed ones an average value of $J_0 = J_{EE} - \delta J_-$, where J_{EE} describes the unstructured network. In the example shown in Results Section, the number of stimuli was p=24 ; the coding level (sparseness) f=0.083; $J_1 = 2.2 J_{EE}$; $J_0 = 0.8 J_{EE}$.

Theoretical estimates of stable rate distributions

We recapitulate briefly the results of the analysis for comparison with simulations. For technical details see (AB96) and Appendix A. The assumptions are:

1. Each neuron receives a large number of spikes in an interval corresponding to its integration time period;
2. The depolarization caused by a single spike is small compared to the threshold;
3. Spike emission times of different neurons are uncorrelated, though their rates may be.

Under those assumptions one can approximate the synaptic input of a neuron by a Gaussian diffusion process, whose mean and variance density can be derived as a function of: the average spike rates of excitatory and inhibitory neurons; the average synaptic efficacies and their variability; the number of afferent excitatory and inhibitory connections; and the neuron's integration time constant.

From these data one deduces the average firing rate of both types of neurons in the local module, as a function of the mean and variance of the input received by the neuron. Since these moments depend in turn on the average rates, the analysis deduces self-consistent equations giving spontaneous self-reproducing average firing rates in the network. (See Appendix A). The stability analysis shows that for realistic rates of spontaneous activity, of the order of 1–5 spikes/second (s^{-1}), to be stable, inhibition has to be stronger than some critical value depending on the ratio of the integration time constants of excitatory and inhibitory neurons. The theory was then extended to the case in which learning occurs in the column. It is found that 'working memory' states, representing the stimuli presented during training to the network, can be stable in absence of the stimulus, provided Hebbian potentiation is strong enough. This simple theory was used to select parameter regions for the simulations. *In all cases where analysis predicted stable dynamics the full simulations were stable and produced the expected average rates.*

The analysis (AB96) was carried out, for simplicity, in a situation in which all neurons receive the same number of connections. It has been extended to the case in which the number of connections varies from neuron to neuron, leading to a wide distribution of firing rates in the network. The equations for both average firing rates and their distribution are presented briefly in Appendix B. They will be described in more detail elsewhere.

Simulation procedure, recordings and display

In the simulation the neural dynamics, Eq. (1), is discretized and the membrane potential at $t+dt$ is given by

$$V_i(t+dt) = (1-\frac{dt}{\tau_\alpha})V_i(t) + \sum_{j\in E}\rho_j(t)J_{ij}S_j(t-\tau_{ij}) - \sum_{j\in I}\rho_j(t)J_{ij}S_j(t-\tau_{ij}) + I_i^{ext}. \quad (4)$$

in which the delays, τ_{ij} are in units of dt, and the external input, I_i^{ext}, is given by Eq. (3). If in $t+dt$ V_i reaches threshold, the neuron emits a spike, as indicated by a spike activation variable $S_i(t)$ (=0,1). The integration step dt was 0.01–0.05ms for both types of neurons, an interval in which on average a neuron receives less than one spike. Lowering dt does not significantly modify the results.

The duration of the runs varied between 200ms (short test runs) and 3 minutes in real neural time, as fixed by setting the absolute refractory period to 2ms.

The simulations of the unstructured network were done as follows:

1. The synaptic structure was set up according to the procedure defined in 'Recurrent synaptic interactions' above.

2. The initial depolarization of the neurons is was drawn from a Gaussian distribution of mean and variance given by assuming that spike emission processes of different neurons are independent Poisson processes with the expected spike rates (AB96 and Appendix A) truncated at the neural threshold θ. We checked that choosing different distributions, (for example all neurons initially at resting potential) leads to the same results, apart from a transient period of about 50ms.

3. At each time step of the simulation the states of all neurons are updated according to Eq. (4), and the following observables are recorded:

- Global activity: total number of spikes in excitatory and inhibitory populations in dt at time t:
$$n_{E,I}(t) = \sum_{i\in E,I} S_i(t); \quad (5)$$

- Single neuron recordings: spike times and depolarization of a subset of randomly chosen neurons (20 excitatory and 10 inhibitory).

At the end of a run we calculate the population averages of the rates, $\nu_{E,I}$, for excitatory and inhibitory networks,

$$\nu_{E,I} = \frac{1}{N_{E,I}T\,dt}\sum_{t=1}^{T} n_{E,I}(t). \quad (6)$$

From the global activity at each time step we compute the *average* cross-correlation (CC) between a pair of neurons one of type α and one of type β:

$$C_{\alpha\beta}(\tau) = \frac{1}{N_\alpha N_\beta(T-\tau)}\sum_{t=1}^{T-\tau} n_\alpha(t)n_\beta(t+\tau) \quad (7)$$

where T is the length of the run in units of dt.

Asymptotically, for large τ, the average CC, Eq. (7), tends to $(\nu_\alpha dt)(\nu_\beta dt)$, with ν_α the mean rate in population α. This background value is used in the following as a normalization, dividing the RHS of Eq. 7.

Single cell recordings are visualized as in physiological experiments – spike times as rasters and depolarizations as continuous curves. From the spike times of cell i we compute its spike rate, ν_i, averaged over an interval T as

$$\nu_i = \frac{1}{Tdt} \sum_t^T S_i(t),$$

which gives the population distribution of rates, and its ISI (inter spike interval) histogram.

From the spike times of a given pair of cells i and j we compute the cross-correlation (CC) of their spike trains, defined as

$$C_{ij}(\tau) = \frac{1}{T-\tau} \sum_{t=1}^{T-\tau} S_i(t) S_j(t+\tau). \tag{8}$$

Averaging C_{ij} over all pairs of cells, one of type α and one of type β, one recovers the average crosscorrelation $C_{\alpha\beta}$, Eq. 7, that is

$$C_{\alpha\beta}(\tau) = \frac{1}{N_\alpha N_\beta} \sum_{i \in \alpha, j \in \beta} C_{ij}(\tau),$$

which was expressed above in terms of the global activities. $C_{\alpha\alpha}$ is the autocorrelation of the global activity.

Results

Spontaneous activity – a typical example

We start with a network of $N = 7{,}500$ neurons, of which $N_E = 6{,}000$ and $N_I = 1{,}500$. The parameters of the network are summarized in Table 1. The connectivity, $c = 0.2$, is random and quenched. The synaptic parameters were determined by theory (AB96) obtain an expected average spontaneous rates of ν_E=13s^{-1} and ν_I=25s^{-1}. The rates are high to compensate for the small size of the network. The statistics of the excitatory afferent input here would be the same as in a network with an average of 8000 recurrent excitatory connections per neuron and an average excitatory rate of 2s^{-1}. In a stationary situation, each excitatory neuron receives about 160 recurrent excitatory spikes, 160 excitatory external spikes, and about 80 inhibitory spikes per τ (=10ms). Each inhibitory neuron receives about half these numbers per τ_I, because of its shorter time constant.

Fig. 1 presents the 'intracellular' recordings of the depolarization of randomly chosen neurons (4 excitatory and 1 inhibitory), and the spike rasters of randomly

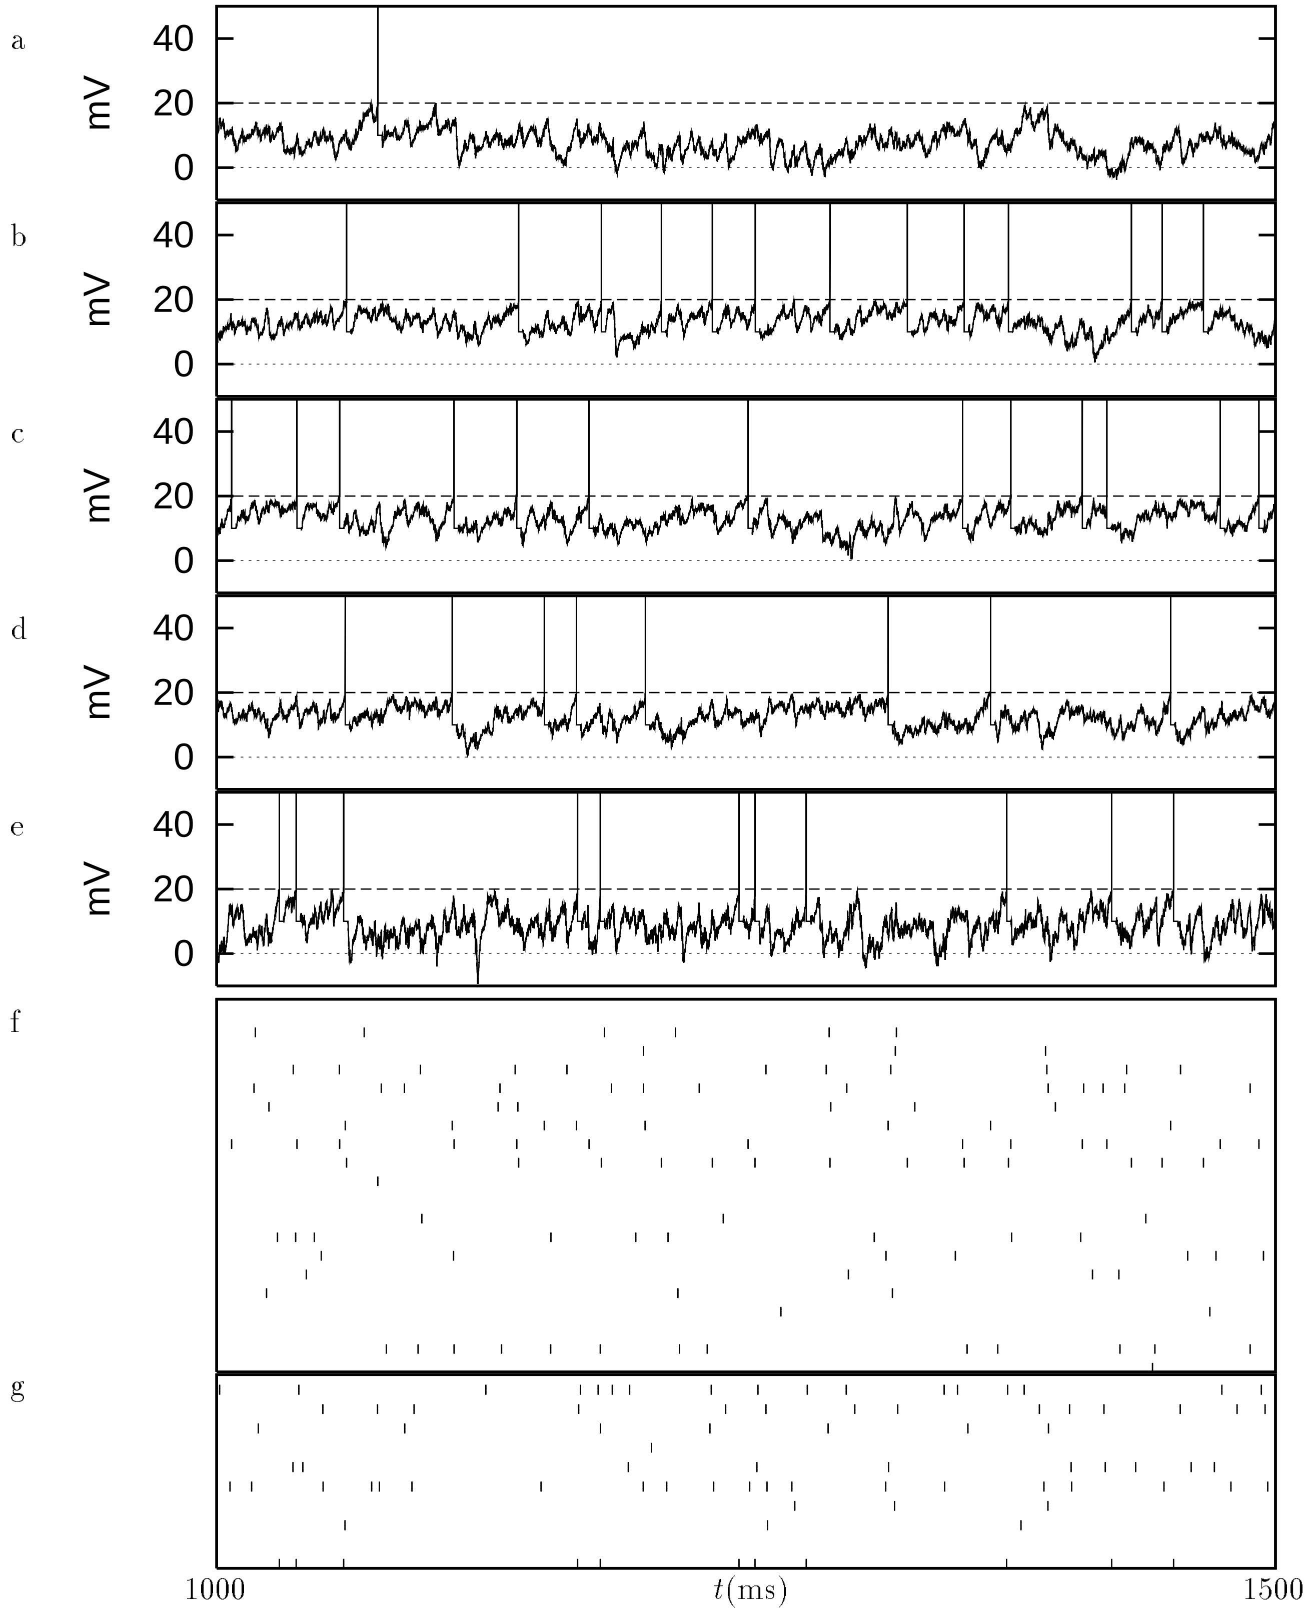

Figure 1: (a)–(e): 'Intracellular recordings' of depolarization in 500ms sub-interval of the run. From top: 4 excitatory neurons 1 inhibitory neuron. Dotted line: Resting potential. Dashed line: Threshold. The size of the depolarization at the spikes is only a visual aid. (f): Spike rasters of 20 excitatory neurons. (g): rasters of 10 inhibitory neurons.

Neuronal parameters	E	I
Number of neurons	6000	1500
τ (ms)	10	5
θ (mV)	20	20
reset potential (mV)	10	10
Observed frequencies (Hz)	13.6	25.5

Synaptic parameters	E→E	I→E	E→I	I→I
Average efficacy (mV)	0.21	0.63	0.35	1.05
Probability of connection	0.2	0.2	0.2	0.2
Average delay (ms)	1	1	1	1
Average number of afferent spikes/τ	163	77	82	38

Table 1: Parameters used in the simulation of Section: Spontaneous activity — a typical example.

selected (20 excitatory and 10 inhibitory) simultaneously recorded neurons in 500ms out of a 100-second run. The horizontal dotted line indicates the resting polarization and the dashed line is the threshold. The neurons are most of the time depolarized and emit spikes due to abrupt fluctuations.

Distribution of firing rates

Fig. 2 is the most remarkable result of this study: the histograms of the measured spike rate distributions in the two populations of neurons (to the left: excitatory; to the right: inhibitory) are superposed on the theoretical predictions. This is not a fit of the results of the measurements (simulations), but the prediction of the theory, which given the parameters of Table 1 gives the rate distribution via Eqs. (19-21) of Appendix B. A χ^2 test shows that the measured rates are highly consistent with the theory.

The distributions have large widths, approximately equal to the average firing rate. It is mainly due to the variability in the number of connections received by the neurons. The effect of this variability, which is relatively low, is 'magnified' by the fact that neurons are most of the time near threshold.

Global activity fluctuations and average cross-correlations

In an unstructured network the total number of excitatory and of inhibitory spikes emitted at any given time, determines the statistics of spike emission in the next moment. Hence, the dynamics of a network can be followed by viewing the time evolution of these two global variables, displayed in Fig. 3 (lower strips). One observes very large, sharp fluctuations which are underlined by comparison with the activity of a population of independent Poisson processes with the same average fre-

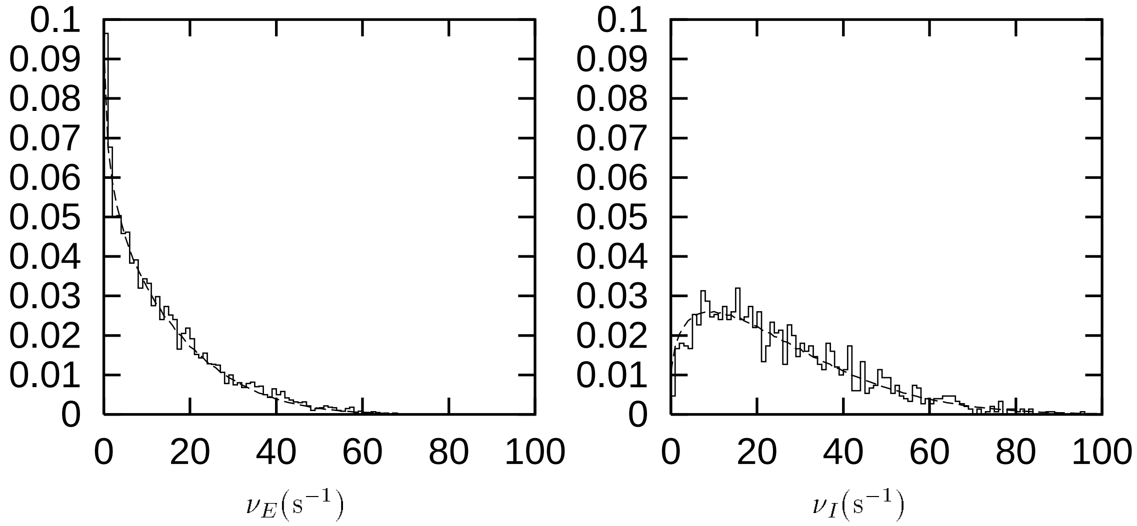

Figure 2: Population distribution histograms of the time-averaged spike rates of all neurons in the network during spontaneous activity. Bin width: $1\mathrm{s}^{-1}$. Left: excitatory neurons. Right: inhibitory neurons. The ordinate is the fraction of cells in the average spike rate bin. The dashed curve is the theoretical prediction.

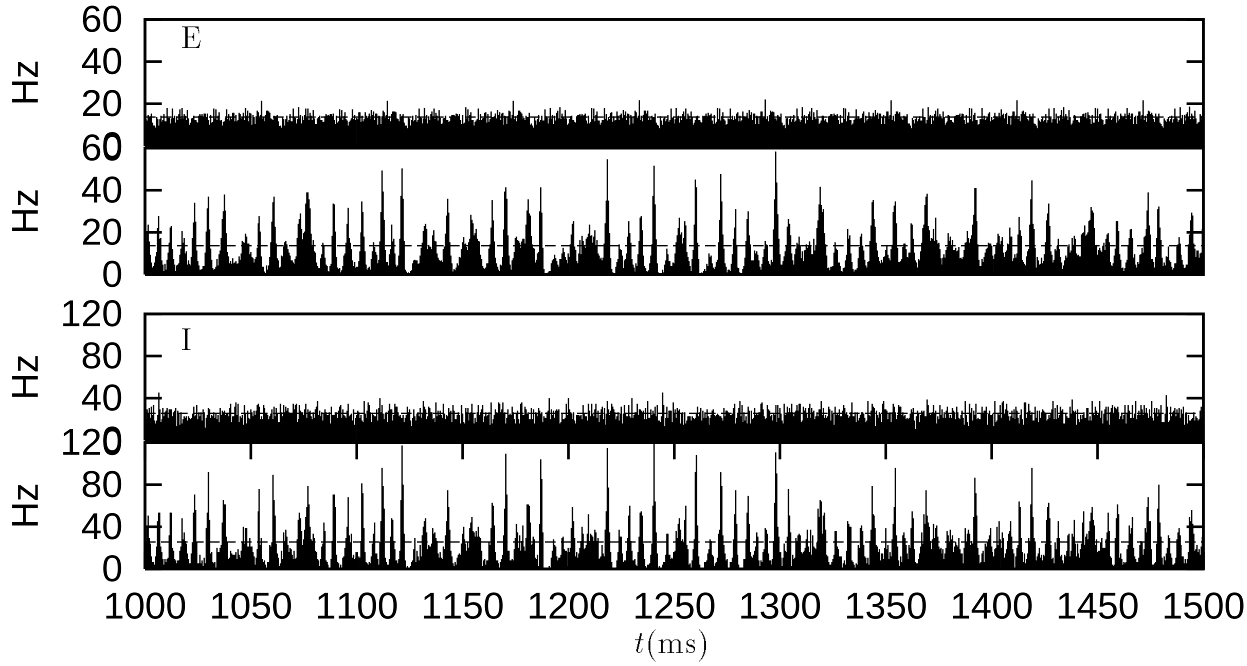

Figure 3: Global activity in the entire excitatory (E) and inhibitory (I) populations during 500ms vs time (lower strip, in each window), in Hz (fraction of spiking neurons in 0.5ms bin / bin size). Note the very sharp and brief fluctuations, to be compared with the relatively smooth behavior of a network of uncorrelated Poisson processes with rates equal to the average rates (horizontal line) in the simulated network (upper strips, same scale).

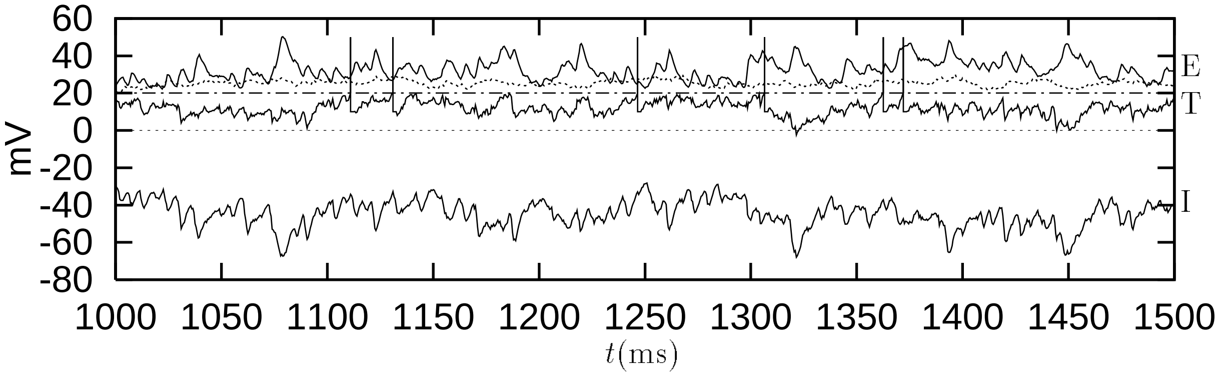

Figure 4: Components of neural polarization of a randomly chosen cell during 500ms: External excitatory (dashed curve), recurrent excitatory (full curve, E) and recurrent inhibitory (full curve, I). Curve T: net membrane depolarization, sum of all three. Spikes are schematized as in Fig. 1. Threshold: dot-dashed line. Resting potential: dotted line. Note the fluctuations of both recurrent excitatory and inhibitory components in correspondence with the fluctuations of the global activity, Fig. 3.

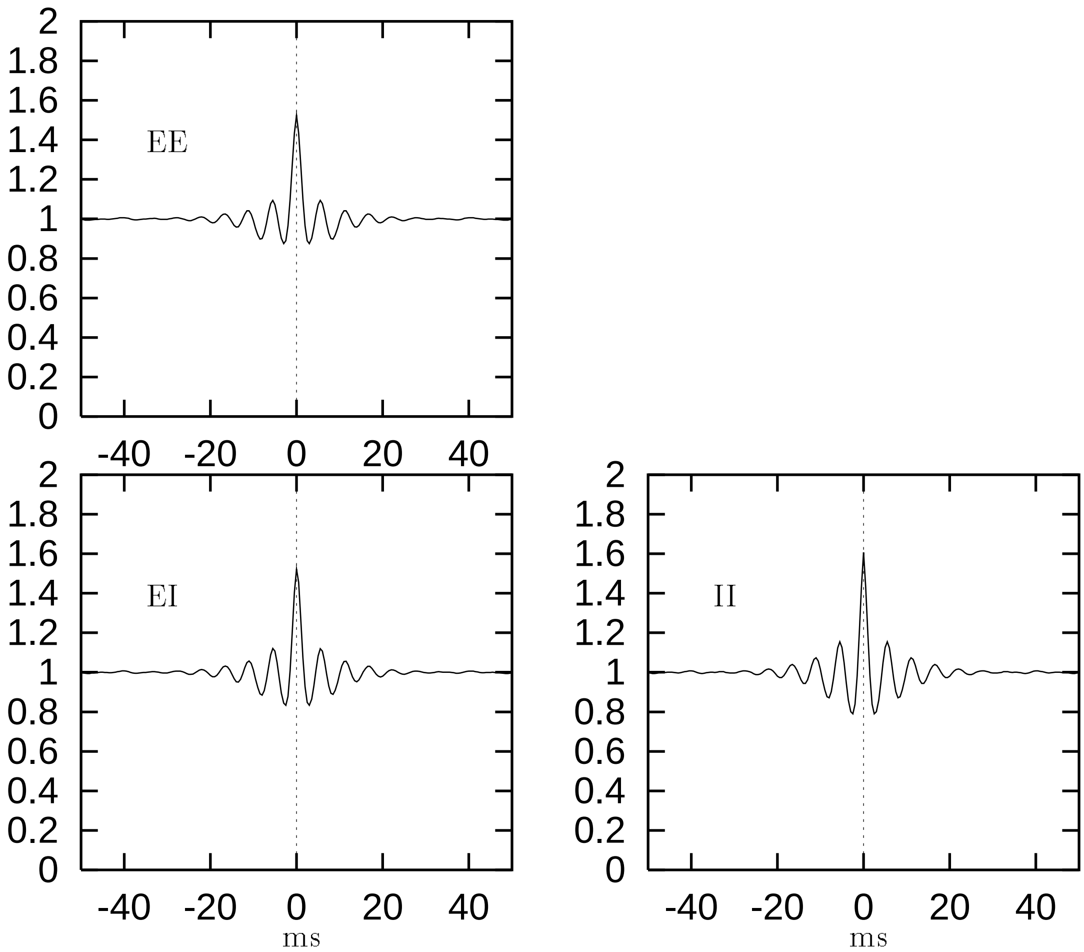

Figure 5: Average CCs between E-E, E-I, and I-I neurons. Note the similarity between the three figures. In all three there is a central peak with the same amplitude relative to the background (between 1.5 and 1.6), and the same damped oscillations with a time constant of about 5ms. The only difference is in the slight asymmetry of the E-I CCs, indicating that inhibitory neurons have a tendency to fire just after excitatory neurons.

quency as in the simulation (upper strip in each window, same scale). The statistical fluctuations, which appear also in the uncoupled Poisson population, are initially amplified by the cooperative dynamics of the network, and only later suppressed by the rising inhibition. If inhibition is significantly weaker, these fluctuations are further amplified and the activity distribution becomes unstable.

The 'surges' in the number of emitted excitatory spikes reach excitatory and inhibitory cells at about the same time and are amplified during the time in which inhibitory neurons integrate the afferent excitation, which is of order 5ms. After this time, extra inhibition is activated, and if strong enough, suppresses the excitation. The inhibitory integration time is roughly the time scale of the sharp fluctuations. These fluctuations are reflected in the collateral contributions (excitatory and inhibitory) to the EPSPs received by neurons in the network. They are shown, for a randomly chosen excitatory neuron, in Fig. 4. The hyperpolarizing inhibitory contribution overpowers the recurrent excitation and both fluctuate with the global activity, Fig. 3.

The activity surges condition the structure of the cross-correlations of spike trains of pairs of neurons. If the number of spiking neurons surges much above average, the probability that any two neurons emit spikes within the short interval of the surge is higher than average. This, in turn, leads to a narrow peak near zero time difference in the CCs. These peaks are observed in all three average CCs in Fig. 5, together with a time scale of a damped oscillation of 5–6ms, that can be interpreted as a degree of synchrony in the neuronal firing.

The structure of the CCs and the implied synchronization appear in contrast with the assumptions of independence underlying theory (AB96). Yet the fact that a simple theory assuming uncorrelated spike emission processes reproduces the rate distribution that well (Fig. 2) is a clear indication that these correlations are effectively small.

Single cell spike statistics

Fig. 6 shows the distribution of inter-spike intervals (ISI) of two randomly chosen neurons. The histograms are similar to the superposed ISI histogram of a Poisson process, apart from the behavior at small inter-spike intervals, which is due to refractory behavior. The variability of ISI intervals, as measured by the ratio of the standard deviation to the mean of the distribution (CV in the figures), is somewhat lower than that of a Poisson process (=1), though some neurons have CV>1. Neurons with low firing rates have CVs larger than 1, while neurons with higher firing rates have CVs typically below unity.

The variability of spike emission can be accounted for by variations in the post-spike reset potential (Tsodyks et al 1994). We find that it is effectively controlled by the relative magnitude of the average depolarization of the neuron to the post-spike reset potential: if the average depolarization is below the post-spike reset potential, variability will be typically higher than 1, and vice versa. Variability decreases as rates increase, at constant V_r, because higher rates are accompanied by higher

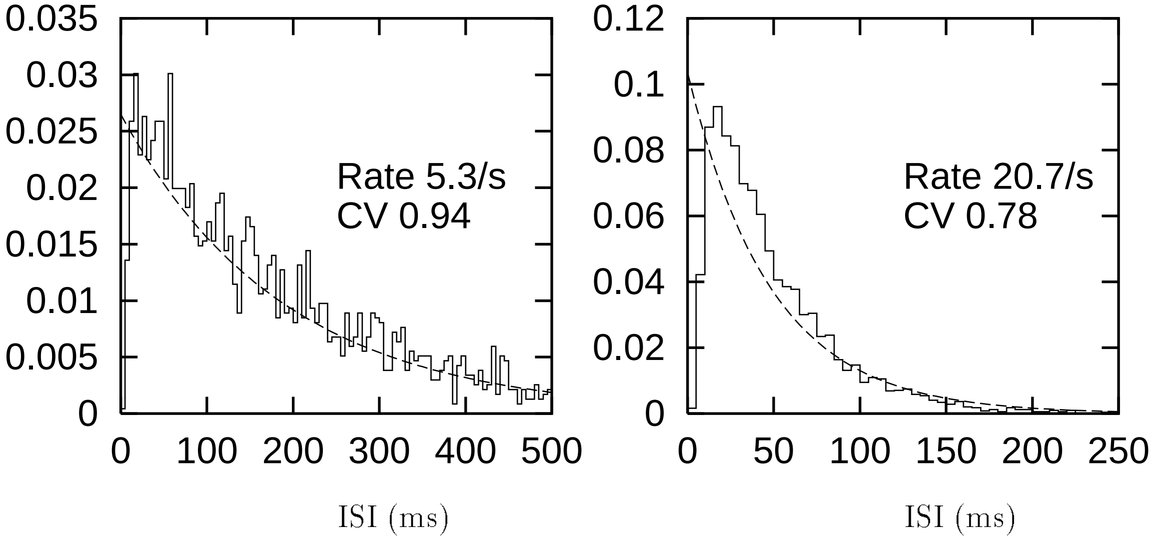

Figure 6: ISI histograms of two randomly chosen E neurons. The average rate and variability (CV coefficient) are indicated in each figure. ISI histograms of a neuron emitting spikes as a Poisson process with the same rate (dashed lines) are superposed for comparison.

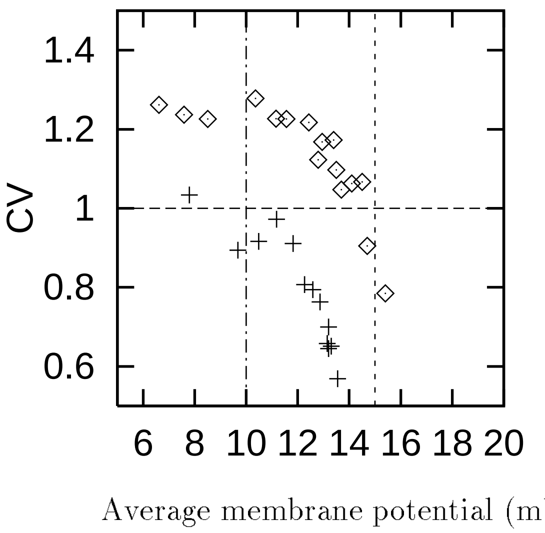

Figure 7: CV of randomly selected excitatory cells (each point one cell) in two networks with different reset potentials vs the average cell membrane potential. $\diamond$: network with V_r=15mV (dashed vertical line). $+$: network with V_r=10mV (dot-dashed line). When the average depolarization of the cell is below the reset potential, the coefficient of variation is typically larger than 1.

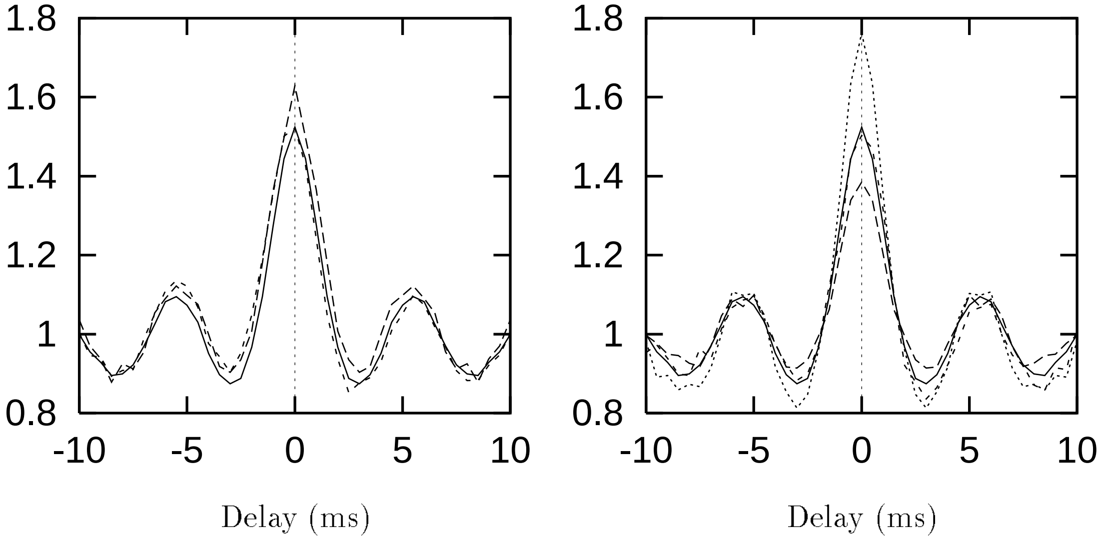

Figure 8: Dependence of average CCs on synaptic connectivity (left) and on rates (right). Left: Full line: average CCs. Long-dashed line: CCs of reciprocally connected neurons. Short-dashed line: CCs of one-way connected neurons. CCs of unconnected cells indistinguishable from average CCs (see text). Right: Full line: average CCs; long-dashed line: CCs of pairs having both high rates ($>30s^{-1}$); short dashed line: CCs of pairs one with low rate ($<10s^{-1}$), the other high rate ($>30s^{-1}$), asymmetric; dotted line: CCs of pairs with both low rates ($<10s^{-1}$).

average depolarization.

We have simulated networks with two values of V_r (10 and 15mV). We measured the dependence of CV of one cell as a function of its average depolarization. The results are summarized in Figure 7. At a given V_r, CV decreases when the average membrane potential increases. It is typically > 1 for cells depolarized, on average, below V_r, and < 1 for cells on average depolarized above V_r.

Cross-correlations

Crosscorrelations are perhaps the most detailed quantity systematically measured in multiple unit recordings. We have studied the dependence of these quantities, in the collective state of spontaneous activity as a function of various parameters.

Effect of individual neuronal parameters.

In a network with parameters as in Table 1, we investigated the dependence of the average CCs in each sub-population on the specifying parameters.

Effect of connectivity: The set of all recorded excitatory neuron pairs was divided into three subsets:
1. Pairs reciprocally connected;
2. Pairs with a one-way connection;
3. Pairs without direct synaptic contact.

The average CC in each class was computed. The results are shown on the left of

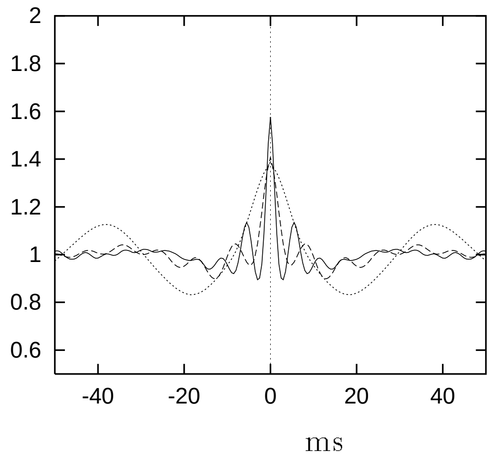

Figure 9: Average crosscorrelations between excitatory neurons, for τ_{delay}=1ms (full curve), 2ms (long-dashed) and 4ms (short-dashed)

Fig. 8. It turns out that there is little difference between cross-correlations of these three types of pairs. Reciprocally connected neurons have slightly higher CCs than average. One-way connected neurons have slightly asymmetric CCs, however the asymmetry is hardly noticeable, and would not be noticeable in any experimental situation. Unconnected pairs have CCs indistinguishable from average CCs, just because the large majority of pairs are unconnected.

Effect of Rates: All excitatory neuron pairs were divided into three subsets:
1. Both cells with low rates ($<10s^{-1}$);
2. Both with high rates ($>30s^{-1}$);
3. One with low rate and one with high rate.

Average CCs were computed in each sub-population, Fig. 8. One finds that neurons with low firing rates are typically more correlated than neurons with high firing rates. There is a slight asymmetry in the CCs of pairs with high-low rates, indicating that high rate neurons have a tendency to fire just before low rate neurons when surges of activity occur. Note that these effects are quite similar to the ones obtained in a network of Hodgkin-Huxley neurons (Hansel and Sompolinsky 1996). All differences are much more pronounced that those due to connectivity.

Effect of spike transmission delays

To test the effect of transmission delays, we have performed simulations with average delays between 1–4ms. Increasing average delays, leads to a minor increase in the average rates. In contrast, delays modify quite significantly the shape of the cross-correlations. Fig. 9 shows the average cross-correlations in the excitatory population, for average delays τ_{delay} =1, 2 and 4ms. Increasing delays decreases the damping factor of the oscillations, with a time constant that increases with the average delay. For larger delays the network needs more time to 'recover' from big surges of activity, leading to a larger time constant of the damped oscillations. The amplitude of the

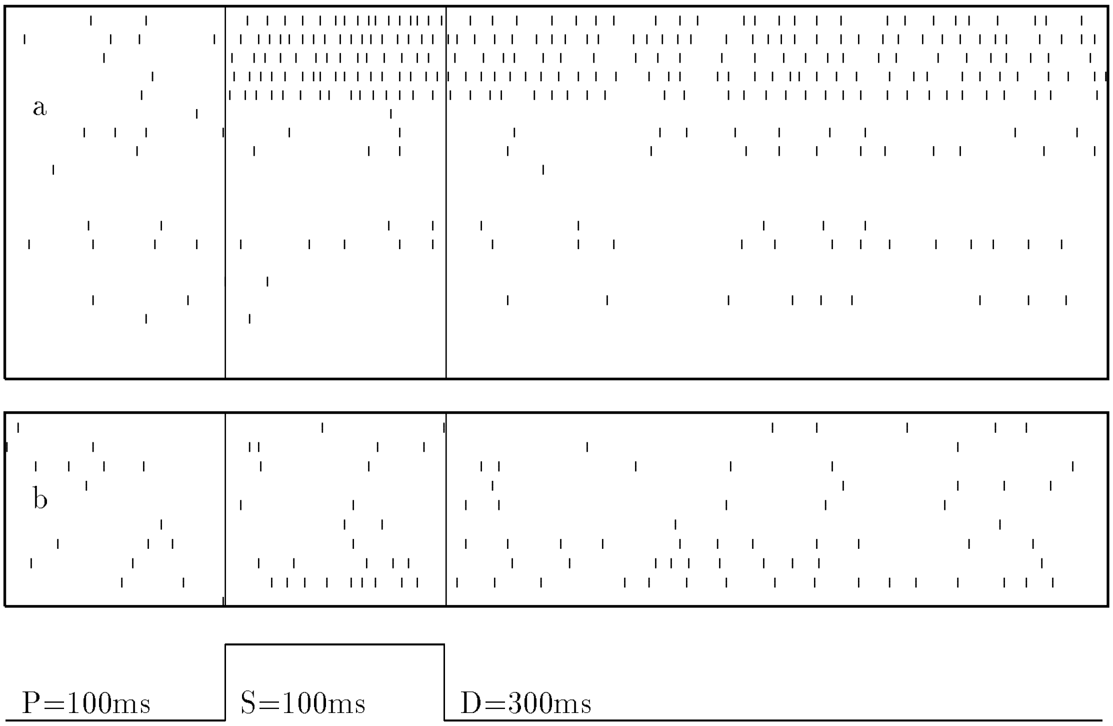

Figure 10: (a): Spike rasters of 20 excitatory neurons. From top to bottom: 5 foreground and 15 background neurons. (b): rasters of 10 inhibitory neurons. The trial protocol is indicated below: P – prestimulus, S – stimulus, D – delay.

central peak somewhat decreases with τ_{delay}.

Other time constants, such as the time constant of conductance changes at excitatory and inhibitory synapses, have a similar effect on the CCs. Simulations performed with finite conductance time constants produce similar 'broadening' of the CCs when the corresponding time constants increase, together with a reduction of the size of the central peak.

Neuronal parameters	E	I
Number of neurons	12000	3000

Synaptic parameters	E→E $J_0 = J_{EE} - \delta J_0$	E→E $J_1 = J_{EE} + \delta J_1$	I→E	E→I	I→I
Average efficacy (mV)	0.10	0.26	0.36	0.21	0.63

Table 2: Parameters used in simulating the structured network (Figs. 10 and 11).

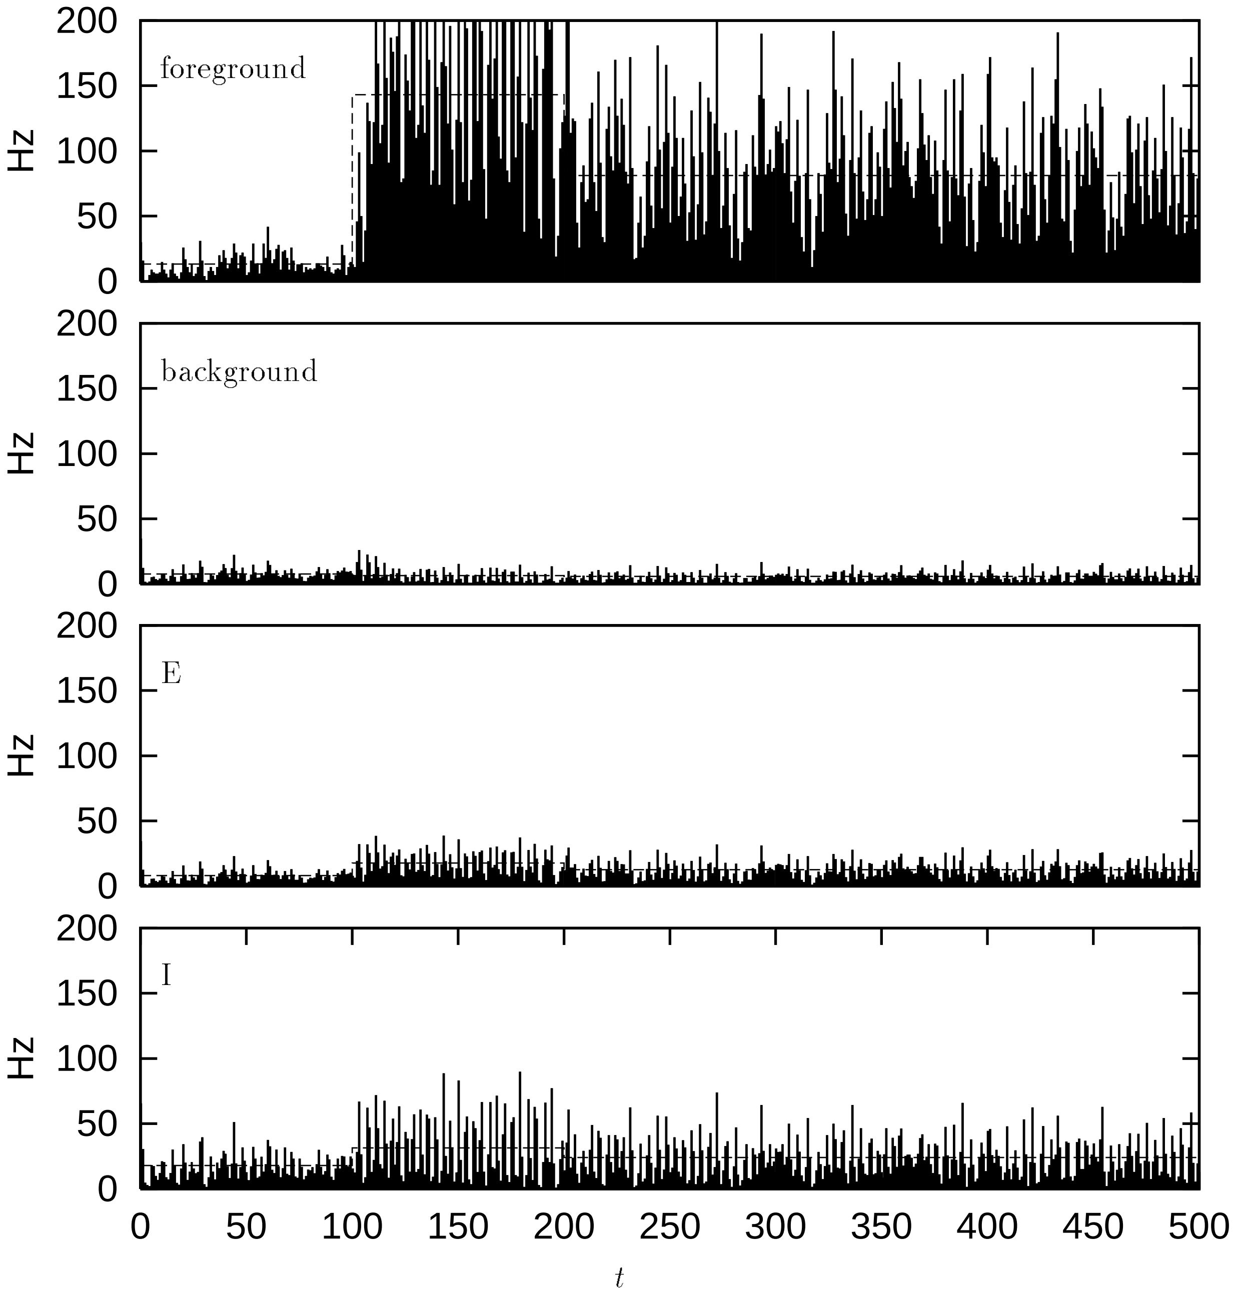

Figure 11: Temporal evolution of global activity (in Hz as in Fig.3) in different neuronal populations in a network stimulated by a familiar stimulus. Windows: foreground of presented stimulus (neurons with 'visual' response); background; all excitatory neurons (E); inhibitory neurons (I). Each window represents the population averaged activity per 0.5ms, together with temporal averages in each of the three intervals of the trial: prestimulus (0-100ms) - stimulus (100-200ms) - poststimulus (200-500ms).

Learning and selective activity

We proceed with an example of the dynamics of the network following learning, to show its operation as an autoassociative memory, with two types of states, a 'spontaneous activity' state and 'working memory' states correlated with learned stimuli. Detailed description and analysis of the dynamics of instructed networks will be given elsewhere. It would make this account too long.

The parameters of the simulated network are given in Table 2. Other parameters are as in Table 1. The structured part of the synaptic matrix corresponds to over-training (repeated presentation) by p=24 stimuli with coding level f =0.08. This creates a synaptic distribution of the form described in Section Methods - The model. The simulation represents the 'testing' phase, following the training. The protocol (bottom trace in Fig. 10) consists of: 1. 100ms prestimulus interval, in which the network is in its spontaneous activity state; 2. 100ms presentation of one of the learned stimuli; 3. Stimulus removed, followed by 'delay' interval of 300ms out of 10s in the actual simulation. We record and display both individual neurons and the evolution of global quantities.

The rasters of 30 neurons (20 excitatory and 10 inhibitory) in the first 500ms of a simulation are shown in Fig. 10. Among the 20 excitatory neurons the top 5 are from the foreground of the presented stimulus (have 'visual response'). The total activity, during the entire protocol, in four neural populations (excitatory foreground; excitatory background; all excitatory; inhibitory) is shown in Figure 11. Before stimulation all neurons in the network are spontaneously active, however this activity is primed by the learning process: neurons in the foreground of *any* of the learned stimuli have a higher activity (about 15s^{-1}) compared to the average spontaneous activity in the excitatory network (about 10s^{-1}). This is seen in Fig. 11, comparing the prestimulus activity in the top window to that in the successive ones. Yet this activity is not selective, i.e. the average rate of neurons selective to any of the learned stimuli is essentially the same.

Upon stimulation[1] by one of the learned stimuli (testing), the stimulated neurons become quickly (in about 5ms) highly active (about 140 spikes per second). Such fast activation is due to the fact that spontaneous activity 'prepares' the foreground neurons to activation, by placing them closer to threshold. Finite synaptic time constants can slow down the neuronal activation (Treves 1993). The other excitatory neurons see their activity lowered from 9s^{-1}to about 7s^{-1}, while inhibition increases from 19 to 32s^{-1}.

After the stimulus is removed the network reaches, again very rapidly, a 'working memory state' correlated with the stimulus, as can be seen in Fig. 11. The foreground neurons stay active in this state at levels of about 80s^{-1}, while the other excitatory neurons stay at about 7s^{-1}. The inhibitory activity is about 25s^{-1}. This state is stable for several seconds, despite the fact that the large fluctuations of the global activity of the unstructured network persist in the memory states.

[1]Since the learning transition probabilities are low, testing can be assumed to take place at constant synaptic matrix.

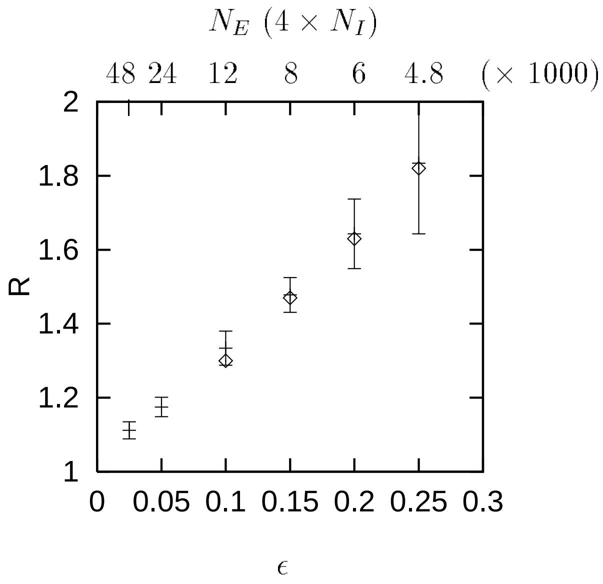

Figure 12: Relative size of the peak in the average CC to its background value $C_{EE}(\tau = 0)/C_{EE}(\tau \to \infty)$ as a function of the transmission probability ϵ keeping ϵN constant. The number of neurons for each value of ϵ are given in upper ordinate. Other parameters are as in Table 1. Error bars are the SD of the average CCs calculated in 1s sub-intervals of the full 10s simulation in the annealed networks. $+$: annealed networks. $\diamond$: quenched networks.

As emphasized in the introduction, all rates (spontaneous, during stimulation, in delay period) are relatively high. This is because the networks are relatively small. In a network of 100,000 neurons we expect the rates to be lower.

Tests related to network size

Connectivities (hence transmission probabilities) and spike rates were increased to reproduce afferent spike statistics of biological situations in networks practicable for detailed simulations. To test the effects of these variables rescaling on various observable quantities while maintaining the afferent statistics, requires increasing the size of the networks. One is severely constrained in the size of networks with fixed synaptic structure that can be simulated. On the other hand much larger 'annealed' networks can be simulated, reaching the size of a cortical column. Wherever possible we have done both. Then we passed to larger sizes with annealed networks. Arguments and experiments justifying this procedure are given in Appendix D.

CCs vs transmission probability

We performed simulations at different values of ϵ, *keeping the average number of spikes arriving at a neuron and all synaptic parameters fixed.* As ϵ decreases, N must be increased to keep ϵN constant. Simulations were run with $\epsilon \in [0.025, 0.25]$. Correspondingly the number of neurons $N \in [4.8 \times 10^4, 4.8 \times 10^3]$, see Fig. 12. If the spike emission processes of all neurons in the network were uncorrelated, the dynamics would not depend on ϵ but only on ϵN, the effective number of connections

N_E	N_I	ν_E	ν_I	J_{EE}	J_{IE}	R_E
6000	1500	10.4	20.8	0.21	0.35	1.83
12000	3000	5.2	10.4	0.19	0.32	2.39

Table 3: Ratio of peak in CC to background for doubled average rate (Right column). Varied parameters of the simulations are left 6 columns. ϵ=0.2. Other parameters are as in Table 1.

received by a neuron. The simulations show that this is not quite the case.

For each set of parameters we performed a 10-second simulation, and calculated the average CCs. The main difference between these CCs is the amplitude of the peak at zero delay, $C_{\alpha\beta}(0)$, relative to the background, i.e. $C_{\alpha\beta}(t)$ at large t. In figure 12 we plot

$$R_E = C_{EE}(\tau = 0)/C_{EE}(\tau = \infty)$$

vs ϵ. It is a measure of the degree of interdependence of the trains of spikes of different neurons in the excitatory network for a given set of synaptic parameters. The corresponding figures for EI or II CCs show essentially the same behavior. Fig. 12 indicates that R_E decreases linearly to 1 as $\epsilon \to 0$.

The average frequencies vary mildly (by less than 10%) with ϵ, for ϵ in the range 0–0.3. This is rather surprising, since as $\epsilon \to 0$, the neurons become uncorrelated, while for finite ϵ correlations are present. In fact, the sharp fluctuations in the global activity, which produce the peak in the cross-correlations, are strongly affected by the variation of ϵ. The surges of activity are followed by relatively quiet intervals, caused by the intervening inhibition, and the net effect on rates turns out to be relatively small. The fact that a rather large peak in the CCs has only small effects on the rates is one of the reasons for the good agreement between simulations and theory.

Simulations with other sets of synaptic values confirm that this phenomenon does not depend on the particular set of parameters: J_{EE} was varied from 0.1mV to 0.4mV. The behavior of the CCs as a function of ϵ remains very similar.

Effect of average firing rates

To check the effect of spike rate scaling on the amplitude on the CCs, simulations were performed varying the *average* firing rates (in contrast to the low rate neurons selected in Section Effect of individual neuronal parameters), keeping the transmission probability and the average numbers of afferent excitatory and inhibitory spikes arriving at a neuron, fixed. This again is done increasing the number of neurons N: keeping νN and ϵ fixed. The parameters of the simulations and the results are shown in Table 3. It shows that the average peak in the CCs *increases* with decreasing average spike rate, much as for different groups of neurons in a given network. Yet theory predicts quite accurately the firing rates in a network in which the peak

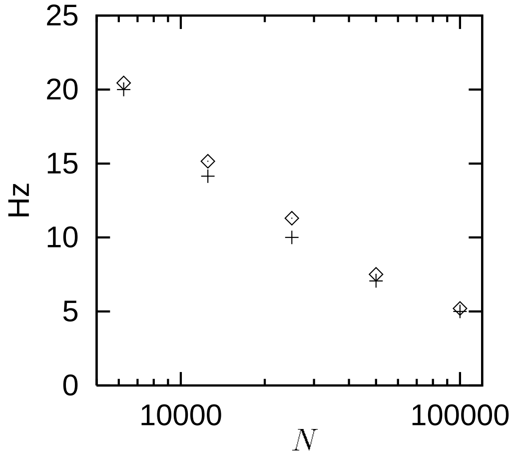

N_E	N_I	ϵ	ν_E	ν_I	J_{EE}	J_{IE}
5000	1250	0.2	20	40	0.16	0.29
10000	2500	0.14	14	28	0.15	0.27
20000	5000	0.1	10	20	0.14	0.25
40000	10000	0.07	7	14	0.14	0.24
80000	20000	0.05	5	10	0.13	0.22

Figure 13: Average excitatory firing rates as a function of N (in log scale). Parameters for each value of N are given in the table. Other parameters are as in Table 1. $\diamond$: simulations. $+$: theory.

in the CCs exceeds two times the background. Thus correlations even at this level have a small influence on firing rates.

Agreement with theory at different N

Fig. 2 demonstrates remarkable agreement of theory (AB96) with experiment (simulation). This agreement is maintained upon varying N from 6250 to 100,000, keeping ν/ϵ fixed, as shown in Fig. 13 (parameters are given in the corresponding table) although the correlation coefficient R_E increases from 1.3 to 1.68, in this range.

Discussion

Simulation as experiment and theory

Networks of IF neurons in realistic conditions exhibit a very rich phenomenology, which shows a remarkable similarity with available neurophysiological data at the level of both mean firing rates and spike statistics beyond average rates. Thus it does not seem necessary to model the individual neuron in great detail to account for currently observed features of cortical spike dynamics *in vivo*. It is very rewarding that a simple theory (AB96) captures quite accurately major aspects of the system: it identifies with precision the regions of the networks dynamic stability; it predicts with no fitting parameters the spike rate distribution; it predicts accurately the level of synaptic potentiation required for selective activity to coexist with spontaneous activity. The simulations confirm the importance of strong inhibitory feedback to control the fluctuations of activity and of the external background currents which 'thermalize' the local module. Lowering the relative inhibitory strength eventually renders the network 'epileptic': all neurons spike at very high rates. Above a critical value of inhibition, low spontaneous activity is stable. The effect is very robust, and *no fine tuning of inhibition vs excitation is required* (AB96). The same robustness is found with respect to the variation of external currents: the spontaneous state is stable to such variations, though average rates vary with the external currents. When these currents alone provoke a depolarization which is significantly lower than the threshold, the network may transit into an almost silent state. All these effects are predicted quantitatively by theory.

Our model departs from other recently studied models at several points: there is no geometrical structure; the dynamics of the neural assembly is studied in the absence of the driving stimulus, though the activity in the network might have been provoked by a stimulus; structured delay activity is studied in conjunction with the stability of the underlying spontaneous activity; it models no predetermined function, but follows the autonomous spike dynamics sustained by the synaptic structure and their development as a consequence of a specific learning dynamics; the synaptic structure is either unformed or structured by a well-defined synaptic dynamics generated by external stimuli; it aims at both qualitative and quantitative phenomena observed in such regions as inferior-temporal cortex; it identifies conditions under which such dynamics can be studied theoretically.

We have found that in networks with a low 'transmission probability', correlations between spike emission processes as expressed by cross-correlations, influence rates weakly. Another indication that these correlations, of experimental magnitudes, are effectively small is the excellent agreement of the rate distribution with theory. We find that *CCs are essentially independent of the synaptic connectivity of the particular pair of neurons*. This finding is quite important for experiments that try to deduce information about the synaptic structure from CCs. CCs depend more strongly on firing rates: neurons with low firing rates have more correlated spike emission processes than neurons with high rates. Delays appear to have a strong influence on the time constants of the damped oscillations in the CCs. We have also found that the dynamics of the network is largely independent of the implementation of synaptic disorder, at equal 'transmission probability'. Networks with fluctuating, unreliable synapses have similar dynamics as networks with quenched, dilute synapses. This could be very useful for the simulation of realistically large cortical modules on modest computers.

Comparison with neurophysiological data

Distribution of firing rates in spontaneous activity. A rather wide distribution of firing rates in spontaneous activity is usually observed in extracellular recordings in cortex (see e.g. Koch and Fuster 1989). In both theory and simulations a wide distribution of rates stems from the variability in the connectivity from neuron to neuron. Other parameters such as the variability of intrinsic neuronal parameters like the integration time constant would increase this variability.

Rates in 'memory states' vs rates in spontaneous activity. We have presented a single example of selective activity following learning. More detailed account of these states of the network will be given separately. Rates were chosen to be high, to counterbalance the average number of synapses per neuron, which is lower than in a cortical module. Also the coding level f in a memory could not be lowered too much, since a critical number of neurons of the order of several hundreds is necessary to sustain a memory state. Already at this preliminary stage the ratio between firing rates in selective and spontaneous activity is rather low, of order 5. This ratio is not very far from what is observed in delay match to sample experiments (see e.g. Nakamura and Kubota 1995, Williams and Goldman-Rakic 1995). To check that memory states with rates and coding levels similar to what is observed in experiments in IT or prefrontal cortex are stable much larger networks would be required. The end product, of course, should be a validation that a simple theory can accurately account for the phenomenology in a wide range of parameters, eliminating the need for future simulation.

Single cell spike statistics. Cortical neurons exhibit a large variability of interspike intervals (see e.g. Softky and Koch 1993). In our simulations a high variability stems from the fact that firing occurs upon unusually high fluctuations of the membrane potential and is similar to a Poisson process. The precise value of the coefficient of variability depends on the relative value of the average depolarization and the post-spike potential.

Cross-correlations. CCs of nearby neurons measured in vivo often exhibit a central peak at zero delay (Kruger and Aiple 1988, Ts'o et al 1986, Gochin et al 1991 Abeles et al 1995). The use of different types of normalizations make it difficult to make quantitative comparisons for the size of the central peak. Where it can be performed, peaks are of the same order of magnitude as in our simulations. An important problem is the fact that in order to estimate reliably CCs of a single pair of low-firing neurons, a very long recording is necessary. Peak sizes 50% above background would not be detected in data like that of (Abeles et al 1995), which is hundreds of seconds long. CCs with damped oscillations have been observed in visual cortex (e.g. Kreiter and Singer 1996) where rates are much higher, and CCs seem stimulus-dependent, which is beyond the present account. Central peaks of experimental CCs often have widths of 10-50ms, larger than CCs shown in Fig. 5. However the width strongly depends on neuronal time constants, as shown in Fig. 9 and on structured activity which will be the subject of a future study.

The rate dependence of CCs appears in agreement with (Abeles et al 1995). They find that the most significant peak is obtained for the lowest firing rates of both neurons. It seems also that the sizes of of the peaks rank according to firing rates, as in the simulations. This is a direct prediction which has not yet been systematically explored experimentally. The same rate dependence results also from the more complex network simulated by Hansel and Sompolinsky (1996).

Comparison with other models of spiking neurons

Single cell spike statistics. Usher et al (1994) have studied a network of IF neurons with a center surround synaptic structure, as a candidate solution to the high variability observed in single cell distributions of interspike intervals (ISI), in a network with high rates, as in primary visual cortex. Tsodyks et al (1994) proposed other solutions to the variability problem. One is related to the possibility that post-spike reset polarization is above the average depolarization level. Our results are quite similar to those of (Tsodyks et al 1994) which was performed on a more detailed neuronal model. However our ISI histograms are quite stereotyped and all look exponential-like, in contrast to those of (Hansel and Sompolinsky 1996). Bimodal ISI histograms characteristic of bursting neurons could probably be obtained were post-spike reset potential random rather than fixed. It may though be that this feature does depend on specific ion-channel dynamics which lead to systematic bursts.

Synchronization. An important question is the definition of a criterion for synchrony in a network state (Ginzburg and Sompolinsky 1994, Hansel and Sompolinsky 1996). They consider the question in the limit in which the number of neurons becomes infinitely large: if the central peak in the CCs disappears into the background in this limit, the network is said to be asynchronous, if it stays finite, it is defined synchronous. In our networks there is no such clear cut distinction. Our network would be considered asynchronous if $N \to \infty$ keeping the number of connections per neuron fixed. It would be mildly synchronous, if the 'transmission probability' is held fixed. Our conclusion is that any network with excitatory and inhibitory neurons with recurrent connections is likely to have a finite degree of correlation in spike emission processes, expressed in the central peak of the CCs. But for realistic levels of transmission probability, these correlations are effectively small.

Acnowledgements

We are indebted to Stefano Fusi for very helpful discussions in the early stages of this work, and to Misha Tsodyks and Yali Amit for comments on the manuscript. DJA's research has been supported in part by a Human Capital and Mobility grant of the EEC. NB acknowledges a fellowship of Programme Cognisciences, CNRS, France, and the hospitality of INFN, Università di Roma, where this work has been started.

Appendix

A. Stable rates in homogeneous unstructured network

In a network of IF neurons the self-reproducing spike rates for excitatory and inhibitory populations are given by the solutions of the coupled equations

$$\nu_E = \phi_E(\mu_E, \sigma_E) \tag{9}$$

$$\nu_I = \phi_I(\mu_I, \sigma_I) \tag{10}$$

in which the means μ_α and SDs σ_α (α=(E,I)) of the synaptic inputs are functions of the rates of the excitatory and the inhibitory neurons, ν_E and ν_I, via

$$\mu_E = \epsilon\tau_E\left(N_E\nu_E J_{EE} - N_I\nu_I J_{EI}\right) + I_E^{ext}, \tag{11}$$

$$\sigma_E^2 = \epsilon\tau_E\left(N_E\nu_E J_{EE}^2 + N_I\nu_I J_{EI}^2\right) + J_{EE}I_E^{ext}, \tag{12}$$

$$\mu_I = \epsilon\tau_I\left(N_E\nu_E J_{IE} - N_I\nu_I J_{II}\right) + I_I^{ext} \tag{13}$$

and

$$\sigma_I^2 = \epsilon\tau_I\left(N_E\nu_E J_{IE}^2 + N_I\nu_I J_{II}^2\right) + J_{IE}I_I^{ext}. \tag{14}$$

The transduction function ϕ_α of α neurons is (Tuckwell 1988, Amit and Tsodyks 1991)

$$\phi_\alpha(\mu_\alpha, \sigma_\alpha) = \left(\tau_{ARP} + \tau_\alpha \int_{\frac{V_r-\mu_\alpha}{\sigma_\alpha}}^{\frac{\theta-\mu_\alpha}{\sigma_\alpha}} du\psi(u)\right)^{-1} \tag{15}$$

with ψ given by

$$\psi(u) = \sqrt{\pi}\exp(u^2)[1 + \mathrm{erf}(u)], \tag{16}$$

and τ_{ARP} is the absolute refractory period.

B. Stable rates in inhomogeneous unstructured networks

When the number of connections received by a neuron varies from neuron to neuron, as is the case when synapses are chosen randomly and independently between each pair of neurons, the (temporally) averaged rates also fluctuate from neuron to neuron. Mean-field theory (AB96) can be extended to cover this situation and produce the full distribution of rates in the population. Here we give a skeleton description of the theory for completeness. Details will be given elsewhere.

The basic idea is that the spike rate of a cell is determined by the mean depolarization and the variance of the **temporal** course of its depolarization, or in a stationary situation, of the **temporal** course of the mean and variance of its afferent current. These are single cell variables. Given the **cell to cell** variability of the connectivity and of the rates, one can estimate the mean and variance of the **population** distribution of afferents over the cells, and hence of the resulting **population** distribution of rates. In a stable situation the resulting distribution should reproduce the assumed one. It must then be checked for stability.

Suppose that the **population** distributions of **temporal** means of the afferents are Gaussian with **population** means for excitatory (inhibitory), E (I), neurons $\mu_{E,I}$ and **population** variances $\Delta\mu_{E,I}^2$. In that case the **population** distributions of the mean depolarizations can be written as:

$$\mu_{E,I}(z) = \mu_{E,I} + z\sqrt{\Delta\mu_{E,I}^2} \tag{17}$$

with z a random Gaussian variable of mean zero and unit variance. For simplicity we neglect the variability in the variances of the afferents. Their effect is relatively weak. Such a distribution of means leads to a **population** distribution of rates via:

$$\nu_E(z) = \phi_E(\mu_E(z), \sigma_E), \quad \nu_I(z) = \phi_I(\mu_I(z), \sigma_I). \tag{18}$$

These induced rate distributions have population means and variances given by the Gaussian averages:

$$\nu_E = \int Dz \nu_E(z), \quad \nu_I = \int Dz \nu_I(z) \tag{19}$$

and

$$\Delta\nu_E^2 = \int Dz \nu_E^2(z) - \nu_E^2, \quad \Delta\nu_I^2 = \int Dz \nu_I^2(z) - \nu_I^2. \tag{20}$$

where $Dz = \exp(-z^2/2)dz/\sqrt{2\pi}$. The **temporal** means $\mu_\alpha(z)$ which enter the computation of the rates, Eqs. 18, are functions of the **population** averages and of the variances of the rates of the excitatory and the inhibitory neurons, via the **population** averages of the depolarization and of their variances, $\mu_{E,I}, \Delta\mu_{E,I}^2, \sigma_{E,I}^2$, which enter Eqs. 17. Those, in turn are given by:

$$\Delta\mu_E^2 = \tau_E^2 \epsilon \left[(1 + \Delta^2 - \epsilon) \left(N_E \nu_E^2 J_{EE}^2 + N_I \nu_I^2 J_{EI}^2 \right) + N_E \Delta\nu_E^2 J_{EE}^2 + N_I \Delta\nu_I^2 J_{EI}^2 \right]$$

and

$$\Delta\mu_I^2 = \tau_I^2 \epsilon \left[(1 + \Delta^2 - \epsilon) \left(N_E \nu_E^2 J_{IE}^2 + N_I \nu_I^2 J_{II}^2 \right) + N_E \Delta\nu_E^2 J_{IE}^2 + N_I \Delta\nu_I^2 J_{II}^2 \right],$$

where Δ measures the spread of the synaptic efficacies and $\mu_{E,I}$ and $\sigma_{E,I}$ are given in terms of the population averaged rates $\nu_{E,I}$ by Eqs. (11–14). Self reproducibility of the activity state of the system is the condition that the mean and variances of the induced rate distribution are equal to those of the inducing distribution.

Given the rates as functions of z, the distribution of rates in both excitatory and inhibitory networks is computed from:

$$P_{E,I}(\nu) = \int Dz \delta\left(\nu - \phi_{E,I}(\mu_{E,I}(z), \sigma_{E,I})\right) \tag{21}$$

C. Synaptic distributions after learning

When p stimuli $\{\eta_i^\mu\}$ are presented many times in a random order, the probability that the synaptic efficacy J_{ij} end in its potentiated state is proportional to q_+ (the probability that upon activation of both neurons transition occurs) and to P_{ij} – the number of stimuli in the training sequence which tend to potentiate it. The probability of it ending depressed is proportional to q_- and to the number of stimuli

D_{ij} which tend to depress it. In terms of the stimuli variables (Section Structured external inputs): P_{ij} and D_{ij} are given as:

$$P_{ij} = \sum_\mu \eta_i^\mu \eta_j^\mu, \quad D_{ij} = \sum_\mu \left[\eta_i^\mu(1-\eta_j^\mu) + \eta_j^\mu(1-\eta_i^\mu)\right],$$

Since stimuli are shown many times, if at least one of the two neurons (i,j) is activated during the learning process, i.e. $P_{ij} + D_{ij} > 0$, the synaptic efficacy will have made at least one transition. Thus the probability that the synapse stays unmodified is zero, while it is potentiated with probability

$$\frac{q_+ P_{ij}}{q_+ P_{ij} + q_- D_{ij}},$$

and depressed with probability

$$\frac{q_- D_{ij}}{q_+ P_{ij} + q_- D_{ij}}.$$

If $P_{ij} + D_{ij} = 0$ the synaptic efficacy is not modified with probability one. In the simulations we used q_+=0.1, q_-=0.025.

D. Comparison of different types of synaptic disorder

Simulations of quenched networks require storing the full synaptic matrix. Memory sets a limit of about 15,000 neurons on a workstation. An 'annealed' network requires no memory for the synaptic matrix, which allows the simulation of networks with up to 10^5 neurons, at which point the time constraint intervenes. Hypotheses related to the size of the network up to biological sizes can be tested on 'annealed' networks. In order to be able to employ this expediency two issues have to be investigated:

1. Annealed networks are inherently uniform, i.e. all neurons are equivalent with respect to the connectivity and transmission of spikes. Their dynamics would naturally be compared to that of 'quenched' networks generated by procedure A;

2. Realistic networks are at least partially 'quenched' and are not uniform. One must compare the dynamics of uniform vs non-uniform quenched networks.

We have performed simulations to compare the spike dynamics of homogeneous networks with 'quenched' and 'annealed' types of disorder. We find that the network dynamics depends mostly on the 'transmission probability' ϵ. Two limits were compared: full quenching (c=ϵ, r=1) and full annealing (c=1, r=ϵ). The delays were chosen accordingly. The transmission probability ϵ varied in the range 0.1–0.3, and the excitatory-to-excitatory average PSP, J_{EE}, in the range 0.1–0.4mV. Annealed and quenched networks of type A produce in all cases very similar results for both firing rates in excitatory and inhibitory populations and for average CCs. In all cases in which it was feasible to perform simulations of a quenched network (i.e. ϵ not too small), results were essentially identical to the annealed case. This is shown by the proximity of the + and the $\diamond$ in Fig. 12. Similar results hold for average firing rates.

Simulations were then performed to test the effect of homogeneity on the dynamics of the network. Synaptic matrices were generated according to procedures A and B of Section Methods. In A all neurons have an equal number of afferent synapses. In B the number of afferent synapses varies from neuron to neuron. The parameters of both networks are given in Table 1. The measured excitatory spike rates in the homogeneous network (about $11\mathrm{s}^{-1}$) are somewhat lower than in the inhomogeneous network (about $13.6\mathrm{s}^{-1}$). Both in close agreement with the theoretical prediction ($10.4\mathrm{s}^{-1}$and $13.1\mathrm{s}^{-1}$, correspondingly). The rates in the homogeneous network are, of course, uniform, in contrast with the wide distributions observed in Fig. 2. These are the main differences observed between the two networks. ISI distributions and CCs are very similar in both cases. The central peak in the cross-correlations is slightly higher in the case of the homogeneous network, which can be attributed to the lower average firing rate. This similarity is observed in all quantities measured.

References

Abeles M, Bergman H, Gat I, Meilijson I, Seidemann E, Tishby N, and Vaadia E 1995, Cortical activity flips among quasi-stationary states, *Proc. Natl. Acad. Sci. USA*, **92** 8616

Amit DJ 1995 The Hebbian paradigm reintegrated: local reverberations as internal representations, *Behavioural and Brain Science* **18** 617

Amit DJ and Brunel N 1995 Learning internal representations in an attractor neural network with analogue neurons, *Network* **6** 359

Amit DJ and Brunel N 1996 Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex, *Cerebral Cortex* **7** No. 2

Amit DJ, Brunel N and Tsodyks MV 1994 Correlations of cortical Hebbian reverberations: experiment *vs* theory, *J. Neurosci.* **14** 6445

Amit DJ and Fusi S 1994 Dynamic learning in neural networks with material sysnapses, *Neural Computation* **6** 957

Amit DJ and Tsodyks MV 1991 Quantitative study of attractor neural network retrieving at low spike rates I: Substrate – spikes, rates and neuronal gain *Network* **2** 259

Amit DJ and Tsodyks MV 1992 Effective neurons and attractor neural networks in cortical environment, *Network* **3** 121

Artola A and Singer W 1993 Long-term depression of excitatory synaptic transmission and its relationship to long-term potentiation *TINS* **16** 480

Ben-Yishai R, Lev Bar-Or R and Sompolinsky H 1995, Theory of orientation tuning in visual cortex *Proc. Natl. Acad. Sci. USA* **92** 3844

Bernander O, Douglas RJ, Martin KAC and Koch C 1991, Synaptic background activity determines spatio-temporal integration in single pyramidal cells *Proc. Natl Acad Sci USA* **88** 1569

Bliss TVP and Collingridge GL 1993 A synaptic model of memory: long-term potentiation in the hippocampus, *Nature* **361** 31

Braitenberg V and Schüz A 1991 *Anatomy of the cortex*, (Springer-Verlag, Berlin).

Fuster JM 1995 *Memory in the cerebral cortex* (MIT Press)

Ginzburg I and Sompolinsky H 1994 Theory of correlations in stochastic neural networks *Phys. Rev. E* **50** 3171

Gochin PM, Miller EK, Gross CG and Gerstein GL 1991 Functional interactions among neurons in inferior temporal cortex of the awake macaque, *Exp. Brain Res.* **84** 505

Goldman-Rakic PS, Funahashi S and Bruce CJ 1990, Neocortical memory circuits, *Cold Spring Harbor Symposia on Quantitative Biology*, volume LV, 1025

Gross CG, Rocha-Miranda CE and Bender DB 1972 Visual properties of neurons in inferotemporal cortex of the macaque, *J. Neurophysiol* **35** 96

Hansel D and Sompolinsky H 1996, Chaos and synchrony in a model of a hypercolumn in visual cortex *J. Computational Neurosci.* **3** 7

Knight BW 1972, Dynamics of encoding in a population of neurons, *J. Gen. Physiol.*, **59** 734

Kreiter AK and Singer W 1996 Stimulus dependent synchronization of neuronal responses in the visual cortex of the awake macaque monkey *J. Neurosci.* **16** 2381

Kruger J and Aiple F 1988 Multi electrode investigation of monkey striate cortex: spike train correlations in the infragranular layers, *J. Neurophysiol.* **80** 798

Koch KW and Fuster JM 1989 Unit activity in monkey parietal cortex related to haptic perception and temporary memory, *Exp. Brain Research*, **76** 292

Komatsu Y, Nakajima S, Toyama K, Fetz E 1988 Intracortical connectivity revealed by spike-triggered averaging in slice preparations of cat visual cortex, *Brain Res.* **442** 359

Lacaille J 1991 Postsynaptic potentials mediated by excitatory and inhibitory amino acids in interneurons of stratum pyramidale of the CA1 region of rat hippocampal slices *in vitro J. Neurophysiol.*, **66** 1441

Liley DTJ, Wright JJ 1994 Intracortical connectivity of pyramidal and stellate cells: estimates of synaptic densities and coupling symmetry *Network*, **5** 175

McCormick D, Connors B, Lighthall J and Prince D 1985 Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons in the neocortex, *J. Neurophysiol.* **54** 782

Miyashita Y 1988 Neuronal correlate of visual associative long-term memory in the primate temporal cortex, *Nature* **335** 817

Miyashita Y 1993 Inferior Temporal Cortex: where visual perception meets memory, *Ann. Rev. Neurosci.* **16** 245

Mountcastle V, Talbot W, Sakata H, Hyvarinen J 1969 Cortical neuronal mechanisms in flutter-vibration studied in unasthetized monkeys. Neuronal periodicity and frequency discrimination. *J. Neurophysiol.* **32** 452

Nakamura K and Kubota K 1995 Mnemonic firing of neurons in the monkey temporal pole during a visual recognition memory task, *J. Neurophysiol.* **74**, 162

Rapp M, Yarom Y and Segev I 1992, The impact of parallel background activity on the cable properties of cerebellar Purkinje cells *Neural Computation* **4** 518

Shadlen MN, Newsome WT 1994 Noise, neural codes and cortical organization *Current opinion in Neurobiol.* **4** 569

Softky WR and Koch C 1993 The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs, *J. Neurosci.* **13** 334

Thomson A, West D and Deuchars J 1993a Single axon IPSPs elicited by morphologically identified interneurones in rat neocortex in vitro, *J. Physiol. (Lond)* **473** 173

Thomson A, Deuchars J, West D 1993b Single axon excitatory postsynaptic potentials in neocortical interneurons exhibit pronounced paired pulse facilitation, *Neuroscience* **54** 347

Treves A 1993 Mean-field analysis of neuronal spike dynamics, *Network* **4** 259

Ts'o DY, Gilbert CD, Wiesel TN 1986, Relationships between horizontal interactions and functional architecture in cat striate cortex as revealed by cross-correlation analysis, *J. Neurosci.* **6** 1160

Tsodyks M, Bell A, Mainen ZF and Sejnowski TJ 1994 Why do cortical neurons spike irregularly? *Soc. Neurosci. Abs.* **20** 1527

Tsodyks M and Sejnowski T 1995 Rapid state switching in balanced cortical network models, *Network* **6**

Tuckwell HC 1988 *Introduction to theoretical Neurobiology* (Cambridge: Cambridge University Press)

Usher M, Stemmler M, Koch C and Olami Z 1994 Network amplification of local fluctuations causes high spike rate variability, fractal firing patterns and oscillatory local field potentials *Neural Computation* **6** 795

Williams GV and Goldman-Rakic PS 1995, Modulation of memory fields by dopamine D1 receptors in prefrontal cortex, *Nature*, **376**, 572

Wilson FAW, Scalaidhe SPO and Goldman-Rakic PS 1993 Dissociation of Object and Spatial Processing Domains in Primate Prefrontal Cortex, *Science*, **260** 1955