Simulation in neurobiology – theory or experiment?[1]

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

The basic dilemma

Neuro-physiology aims at deciphering the behaviourally relevant dynamics of large assemblies of neurons. In experiments one either measures the behaviour of small samples of neurons, as in electrode recordings, or of averaged signals of large numbers of neurons, as in EEG, MRI etc. Phenomena observed in electrode sampling represent, necessarily, large scale features, or they would not come up in the sampling. Signal averages (over tens of cortical mm^2, several million cells) are of interest if they represent the concerted dynamics of neurons and synapses. One implication is that some notion connecting the presumed underlying structure (of neurons and synapses) to the large scale dynamics is required, which is a theoretical framework.

The actual details of the underlying system are not fully clarified not even the level of detail actually required for the production of the sampled, or averaged phenomena. One way (perhaps the only way) to proceed is to set up models for the system, simulate them and compare the behaviour of the simulated system on the one hand to an accompanying theory, whose role it is to identify the relevant global degrees of freedom, and on the other hand to sample recordings from a brain.

Large scale simulation is becoming an ever more prominent activity in studying neural systems. It goes all the way from detailed, small scale description of a single cell – as in cable theory models [1, 2, 3], to large scale networks of simplified neurons [4, 5, 6] exhibiting various types of collective dynamics. And systems in between, which mix more complex ionic, neuro-transmitter and neural structure with large scale features[7]. Clearly, any such simulation implies a model. One could "run" the simulation and observe the behaviour of the system under consideration. But this is more like an experimental situation than like a theoretical one. If the simulated system is complex enough, generic statements about its product dynamics would be almost as gratifying and surprising as about an experiment. Moreover, to monitor the system's progress one would have to define and sharpen tools, much like in the experimental situation, since the system generates an enormous amount of noisy data. Thus it appears that the simulation hangs somewhere between the theoretical and the experimental.

The matter may be sharpened further: One may be given the design of a part of the brain, consisting of operational features of the elements: neurons, synapses etc., as well as

[1] A color version of the ms can be found on URL: www.fiz.huji.ac.il/staff/acc/faculty/damita

a scheme of connectivity. The design and the parameters may seem plausible to a neuro-biologist. If what is interesting about the brain must look beyond single cells, one would proceed to explore the operational properties of the assembled system, either as a culture, or in a cortical environment. Testing the collective properties of the neural assembly can serve two ends. First, to test the plausibility of the provided schematics. Second, the collective properties may serve as predictions about the biological system. Such a test can be affected either by a hardware emulation of the discovered recipes, or by computer simulation. We set aside the hardware option for shortage of space, though the valid, specific dimension it represents is well worth a reflection (see e.g. Badoni et al 1994[8] and Van der Velde 1995[9]).

The above considerations are in a context in which **the assembly is considered a mechanical system decoupled from psychology or behaviour**. This is a fundamental commitment: a neuron, a synapse, a neural assembly are first and foremost mechanical (biophysical, biochemical) systems. Behaviour is considered the consequence of an interpretation of some aspects of the dynamics of the corresponding system, in its interaction with other systems (cortical or extra-cortical) employed in the execution of a given computational (psychophysical) task.

Hence to the biologist the simulation may appear as a theoretical synthesis of schematizations. But it should also be experimented with. After all, it is not that neuro-physiological experiment presupposes very detailed characteristics of neurons beyond spike emission. The neuro-physiologist studying the cortex of a performing mammal approaches a system of very many neurons, exposed to stimuli, and tests its response in different situations, without any bias concerning the detailed characteristics of the component neurons. Though he may have such a bias, it does not condition the experimental situation and at best partakes in an explanatory stage. Moreover, the responses of the system tested by the neuro-biologist are often decoupled from behaviour. A perfect case in point are the Hubel-Wiesel experiments. A similar attitude guides delayed-match-to-sample experiments[10].

So to a biologist, the simulated network would be a theoretical construct but also a (dry) empirical specimen on which experiments can be performed to characterize its responses, which are not apriori better known than for the brain. For a theoretician the simulation is an experimental system with complex behaviour, ensuing from the complex assemblage of a very large number of non-linear elements. His task is to construct a concise theoretical framework that can predict the collective characteristics of the network. But it is also a theoretical object inasmuch as it provides predictions for the biological system it schematizes, wherever the analysis of the simulated network remains powerless.

In the following I describe an **example** of the proposed methodology. The context is a specific (though complex) network of spiking elements, a candidate model of a cortical module. The subject is spontaneous activity and its characteristics, a rather rich and little studied phenomenon. Single and multiple unit recording data from the simulation is confronted with experimental data, mostly to give the dynamics of the simulation credibility. To consider it as predictive of the experiment would require an agreement on the constitutive parameters. The fact that many dynamical properties of this network can be studied in detail theoretically, represents the explanatory potential of such simulated models. Wherever simulation results are not yet captured by theory, the simulation can be used to exclude certain explanations of experiment and to predict others, because of the total flexibility it enjoys in terms of varying parameters and selecting what to record.

The structureless cortical component

Suppose, in the spirit outlined above, that neurons are simple point integrators of their afferent currents that emit a spike when the integrated level reaches a threshold, followed by an absolute refractory period and a resetting of the level (an integrate-and-fire IF device[11, 12]). These excitatory and inhibitory elements are then put together in large numbers, $10^3 - 10^4$ in a 4:1 proportion excitatory:inhibitory, (not quite the 10^5 in a cortical column, suggested by anatomy). Connectivity is taken, putatively, totally random. That implies: 0. A neuron of every type can synapse on a neuron of any other type; 1. A random number of direct synaptic neighbours to each cell, constrained by a mean connectivity (probability of direct synaptic connection) given by anatomy (about 10% in cortex); 2. Random distribution of synaptic efficacies (excitatory and inhibitory) centered on some plausible mean for each population of synapses; 3. Random delays for spike transmission associated with each synapse and mimicking the finite extent of dendrites.

It would be unnatural to expect such a structureless network to produce anything beyond spontaneous activity. But even at the level of spontaneous activity neuro-physiological phenomenology is rich. It ranges over: spike rate distributions in different classes of neurons; inter-spike-interval distributions on single cells; cross-correlations (CCs) of spike emission times of pairs of different types of cells and of their dependence on various parameters of the network, etc... . One could then proceed to examine how such a network would behave if exposed to learning of external, structured stimuli.

To complete the description of this sample system we note that anatomy indicates that while at the level of a column connectivity is high, each neuron has on its dendrites about as many synapses from remote neurons as from local (collateral) ones. The remote synapses are all excitatory, since typically, in mammalian cortex, only excitatory neurons have long axons. Long-range inhibition is expressed as long-range excitation of local inter-neurons. Again, putatively, we assume that the spike rates carried by the remote neurons are, on average, the same as those of excitatory neurons in the column. Simple theoretical arguments[15] indicate that stability requires a lower limit on the level of local inhibition, it should at least cancel the collateral (local) excitation (see also [16]). With 4 times less inhibitory neurons, the product of their average rates times their average efficacy should be at least 4 times that of the excitatory value.

The meaning and the intuitions behind this constraint are, briefly, that stability is the tendency of the network to suppress fluctuations in the spontaneous rate and to restore the stable value. The danger is always that a fluctuation will feedback positively and drive the rate, once away from its stationary value, either to very high or to very low rates. The inhibitory dominance condition ensures that when the rates go above the stationary ones, inhibition grows faster than excitation and reduces that rate. When the rate fluctuates downward, the balance turns the other way and the rate increases. For detail see e.g. ref. [15]

Note that such a model would be inadequate for describing a neural circuit operating with a small number of neurons, such as central pattern generators in invertebrates. Most probably the functioning of such systems depends on detailed properties of neurons, on specific connectivity arrangements and timing delays of signal transmission. Such a system would not be expected to maintain its own spontaneous activity, unless again composed of autonomously spiking neurons. Optimism in the possibility of accounting for the dynamics in

Neuronal parameters	E	I
Number of neurons	6000	1500
IF time constant (ms)	10	5
Threshold (mV)	20	20
Reset potential (mV)	10	10

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

Table 1: Parameters specifying the network in its unstructured phase. Synaptic efficacies are peak depolarization at the soma per spike. Efficacy distributions have a width 10% their magnitude; connectivity is a binomial process; delays are distributed uniformly in (0–2)ms.

mammalian cortex must be based on a duality between circuit size and required complexity of elements and composition.

Multi-unit recordings and experiment

The anatomy of the system is schematically depicted in Figure 1. Its constitutive parameters are given in Table 1 (See e.g. [5]). With all the parameters set, a simulation starts from some random distribution of depolarizations, and then the network is left to its autonomous dynamics. How should one classify the dynamics of the network? The question is as baffling as in the experimental situation, despite the fact that one knows all there is to know of the constitution of the system and there is full access to all the dynamic information. Recordings can be effected simultaneously on any number of neurons. Figure 2 exhibits several 6.5-second long (in actual neural time) spike rasters and average spike rate histograms of one "average" excitatory neuron (a) and one "average" inhibitory neuron (b), before during and after the presentation of a stimulus (peri-stimulus PST). Figure 3 presents the distribution of average spike rates in the excitatory and inhibitory populations.

Note how wide the rate distributions are. The origin of this width can be traced to the variability in the number of synaptic contacts due to the randomness in the connectivity. Yet, given the high number of synapses per neuron, the relative variability (ratio to mean) in the connectivity is rather small. This small variability is amplified very significantly by the threshold dynamics of the neural elements.

Figure 4 presents the the auto-correlations (AC) of the spike emission times of one high and one low rate neuron, averaged over 5 trials of 5 second activity. Each window is a superposition of the AC of one neuron from the simulation (yellow) and one (with similar rate) recorded in IT cortex. What is rather striking is the similarity of the experimental and simulated distributions – both are well approximated by a Poisson process, except near spike emission. At low rates, the variability so characteristic of the living situation is a natural feature of the dynamics of the simple model. Consequently, whenever the simulation differs significantly from the recording, the origin of the difference can first be searched in the simulated system, which is so much more accessible. Potential answers can then be examined as candidates for explanations for the behaviour of the biological system. Finally, in Figure 5, we present cross-correlations of spike emission times in 5 seconds of two pairs of excitatory, simulated neurons: one pair with high rates, and one with lower rates. The simulation makes a clear prediction: the size of the peak in the CC, relative to background, decreases as the rates of the two neurons increase.

It is a challenge of considerable value to suggest criteria that distinguish these data from

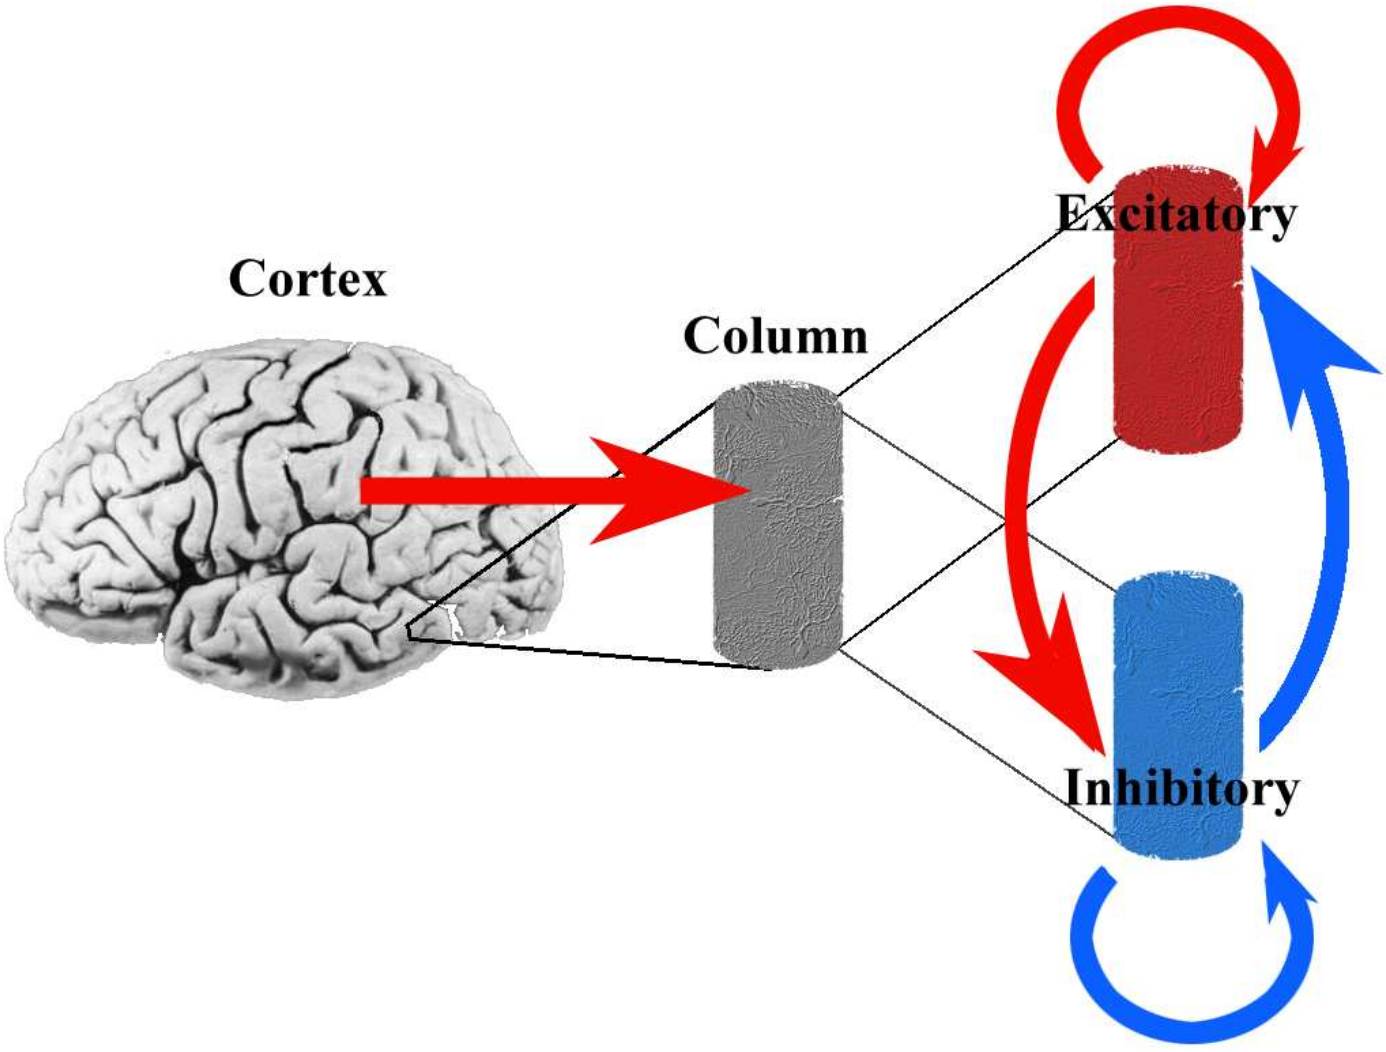

Figure 1: The connectivity scheme of the simulated network. Local excitatory and inhibitory neurons are part of the same local column. Their separation in the figure is intended for clarification of the connectivity. Each connectivity arrow represents a random distribution of connections received from a random sample of presynaptic neurons. Within each connection group the synaptic efficacies are distributed at random and a random transmission delay is associated with each synapse. Only the collateral excitatory-excitatory connections are candidates for learning.

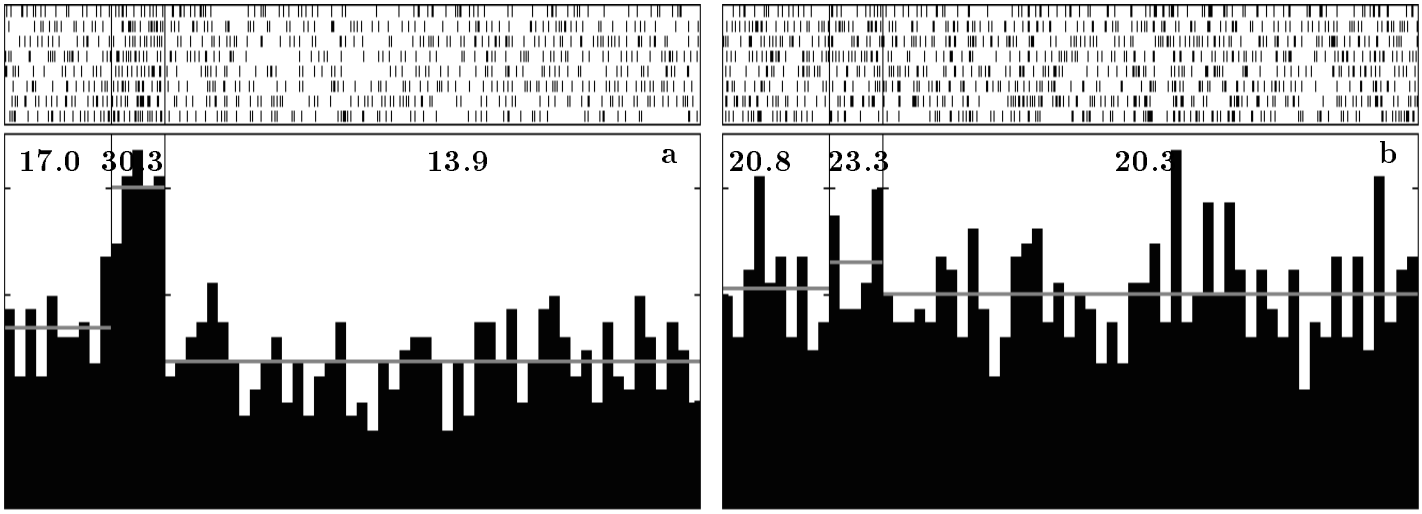

Figure 2: Rasters and PST histograms of spontaneous activity of one excitatory (a) and one inhibitory (b) neuron in 6.5 seconds of simulation of the network specified by Figure 1 and Table 1. First interval 1-second prestimulus; followed by 0.5-second stimulus presentation and 5-second post stimulus activity. Repeated spike rasters represent recording from the same cell in 8 repetitions of the simulation with different randomized sequences. The histogram represents the average rate (spikes/sec) of the 8 runs in bins of 100ms. Horizontal lines in the histograms are temporal averages of the rate in each interval, indicated also numerically. Average rate in each population: 13.6Hz (excitatory), 25.5Hz (inhibitory). The two neurons have been chosen to have rates near their population average.

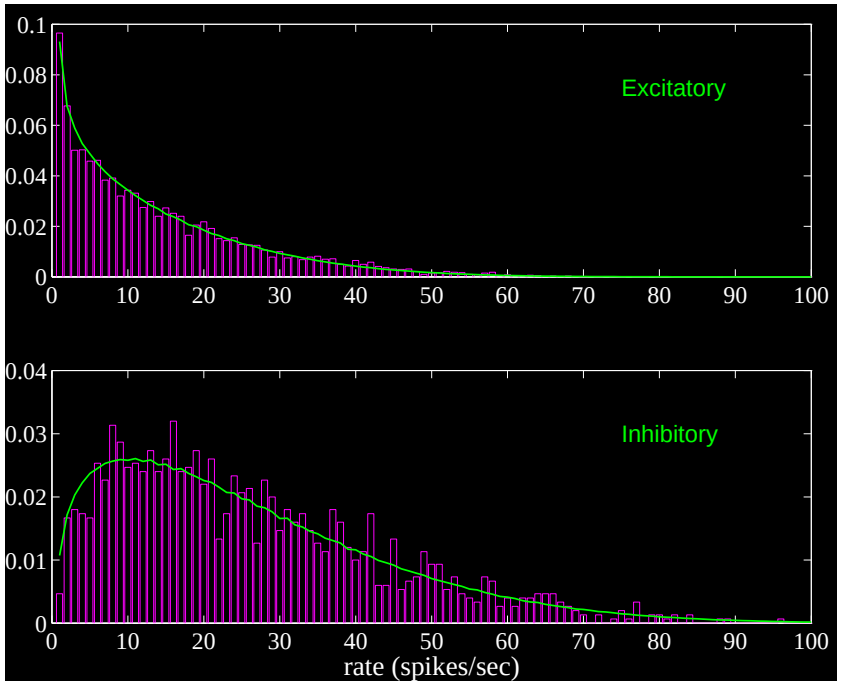

Figure 3: Population distribution of time averaged spike rates, of excitatory and inhibitory neurons, in a structureless network (Figure 1 and Table 1) in a single simulation. The average rate in each distribution is indicated in the figure. Continuous curves are theoretical prediction with the same parameters. Bin width: $1s^{-1}$.

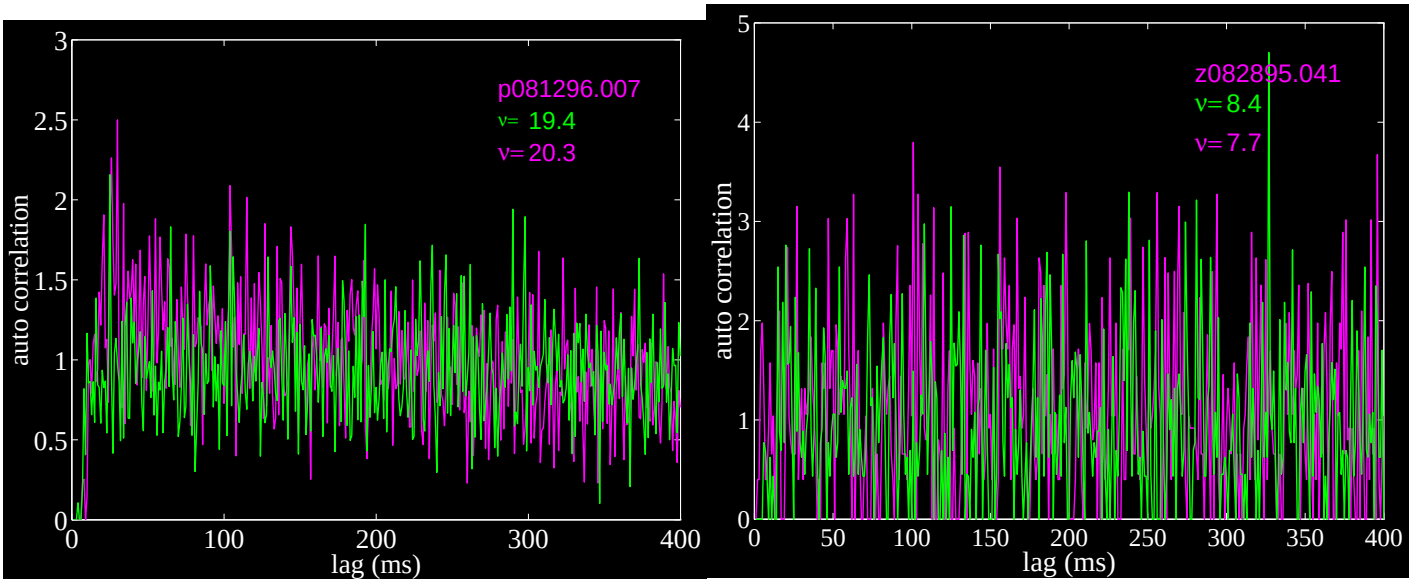

Figure 4: Auto-correlations of spike emission times of two neurons from the simulation and two neurons from recording in infero-temporal cortex of behaving monkey. Each auto-correlation is averaged over 5 repetitions of 5-second activity of the same cell following the presentation and removal of the same stimulus, in cases in which there was no selective delay activity. Left: High rate cells, Right: Low rate cells. Rates are indicated in the figures. Blue: experimental, green: simulation.

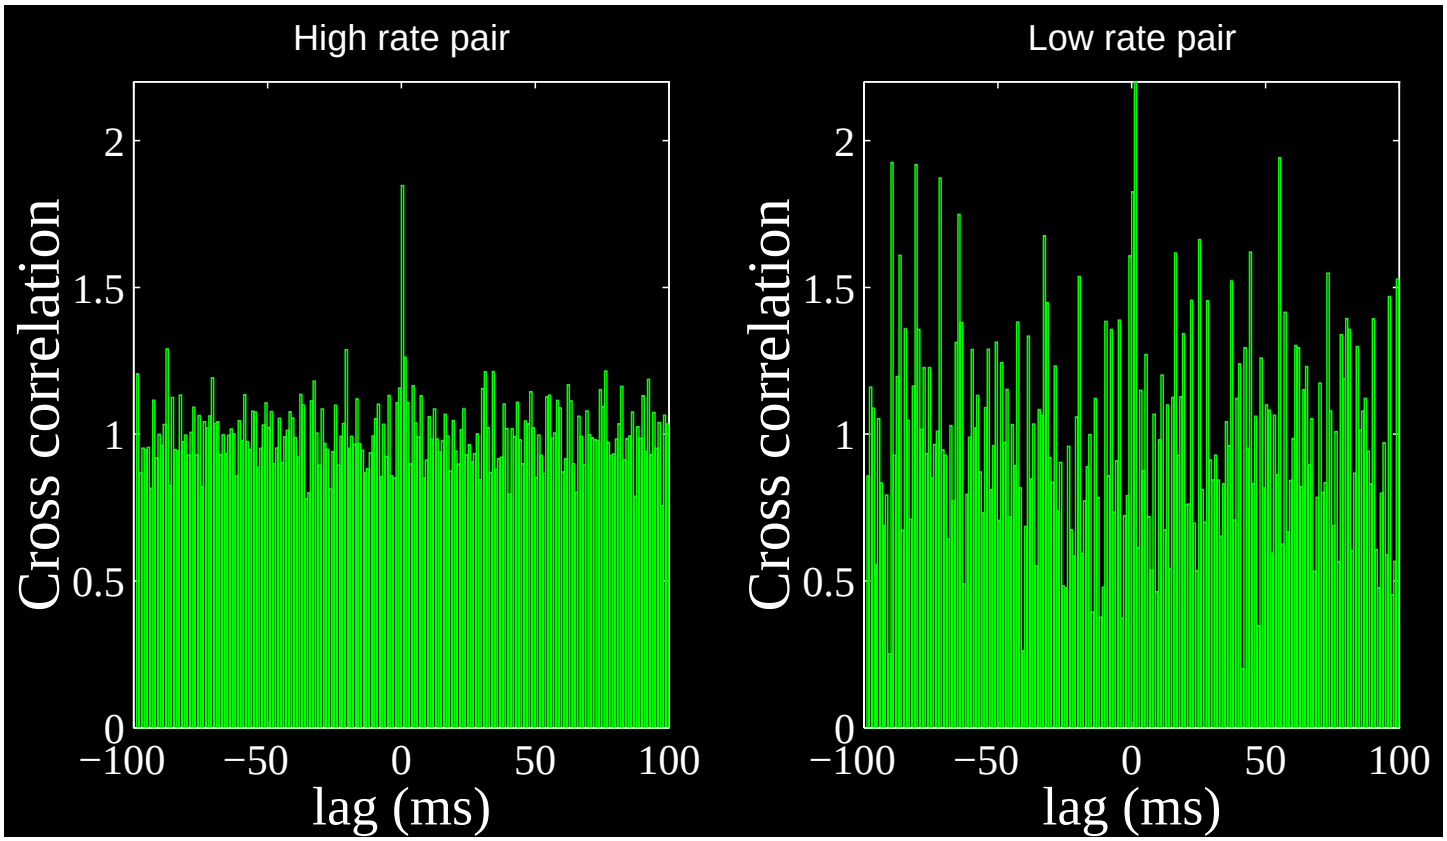

Figure 5: Cross-correlations (raw) of spike emission times of two neuron pairs from the simulation. Left: high rate (about 35Hz) Right: low rate (about 10Hz). Each cross-correlation is averaged over 16 repetitions of 5-second activity of the same pair of cells, following the presentation and removal of the stimulus, with no selective delay activity.

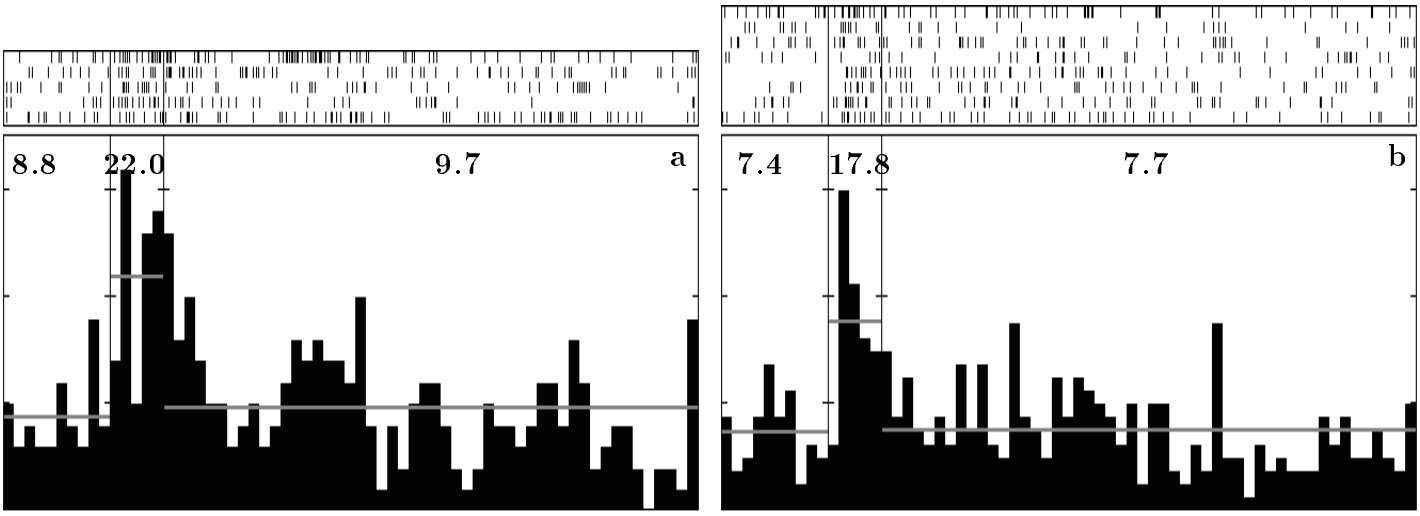

Figure 6: Rasters and PST histograms of spontaneous activity of two cells in inferotemporal cortex of a monkey performing a delayed match-to-sample task[13]. The protocol and notation as in Figure 2. The activity is identified as spontaneous by the absence of selective delay activity (see below). The experiment was not tuned to distinguish excitatory from inhibitory spikes.

empirical neurophysiology in the cortex of a performing mammal. In Figure 6 we show spike rasters and PST histograms; in Figure 7 population spike rate distribution; and in Figure 8 the cross-correlation of a pair of neurons, to underline the challenge. One may try to confront it by pointing out that the experimental PST data (Figure 6) is less regular than the simulated data (Figure 2). This may be a real effect, or it may reflect the simple fact that in the simulation rates are higher than in experiment, to compensate for lower number of neurons. That regularity follows the rates can be seen in Figure 4. It may also be suggested that the experimental CCs (Figure 8) are more sturdy and may express somewhat different time constants than those of Figure 5, either in their oscillations and/or in their decay. The first may again be related to the difference in rates, which this time are higher in experiment, and to the twice higher trial statistics in the experiment. As far as the time constants of the CCs are concerned, ref. [5] indicates that they may be expressing facts about underlying time constants in the network, well under control of the simulation and hopefully soon of theory.

Mean-field theory and simulation

The behaviour of the simulated network is accessible to a mean-field analysis, due to the huge number of connections received by each neuron. The underlying assumptions are that the relevant degrees of freedom are average spike rates; that times of spike emission processes of different neurons are not correlated and that it takes many incoming spikes to reach threshold (see e.g. [15]). In Figure 3 the continuous curves are theoretic calculations of the spike rate distributions in spontaneous activity. The analysis[5] is based on parameters *identical* to those used in the simulation. There is no curve fitting. The same relatively simple theory produces the parameter regimes in which the dynamics would be stable, as well as the variation of the distributions as the parameters are varied. Note that here the simulation, or the corresponding theory, offer a real bonus over the experiment: the freedom to monitor the

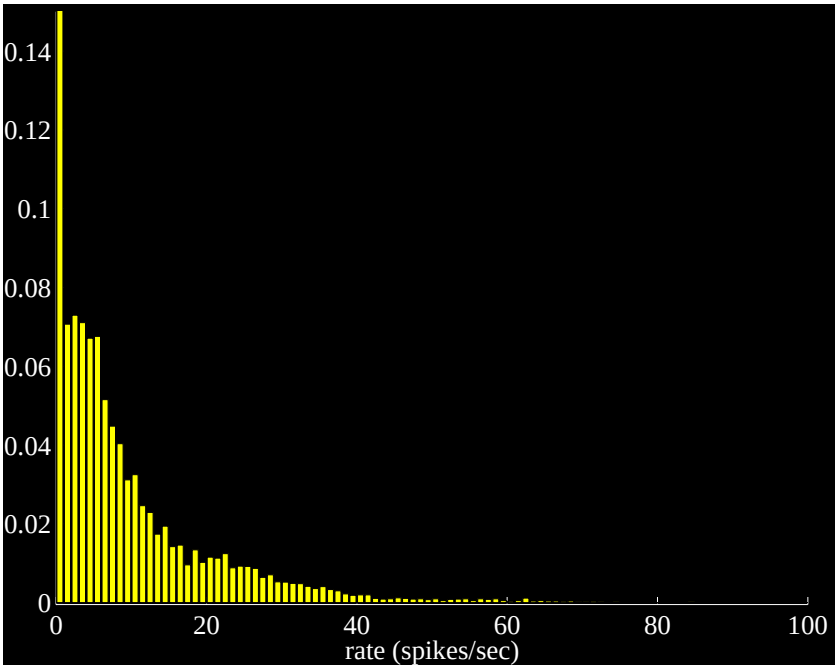

Figure 7: Population distribution of time averaged spike rates, measured in 148 cells in infero-temporal cortex of a monkey performing delayed match-to-sample task[14]. The rates are in a 1-sec prestimulus interval. In the most cases the prestimulus activity is spontaneous. In a few it may be propagated delay activity (see e.g. [13]. Bin width: $1s^{-1}$.

effect of parameter variation. As was mentioned above, the theory describes the dynamics of the network in terms of average spike rates. It is found that whatever can be captured by those, is well captured by the theory, in the sense that it produces a quantitative account of the measurements in the simulation. Questions related to the dependence of rates and to single cell spike statistics on e.g. synaptic strengths, connectivities, neural and synaptic integration time constants, amplitudes of external afferents etc. can be addressed to the simulated 'preparation', and in turn, more generically and more economically to the theory.

The simulations provide predictions for detailed features such as CCs, response times to external stimuli and more. Those are not yet reliably captured by theory. The main limitations of the theoretical analysis are the hypotheses that spike emission times of different cells are independent and that in a given state of the network, spike emission is a stationary process. One possible test of these hypotheses is the measurement of the CCs by 'multi-unit' recordings in the simulations. Were the hypothesis verified, the CCs would be flat, as they are at large values of the spike time-difference. The peaks at small time differences in Figures 5 (simulation) as well as in Figure 8 (experiment) bely the joint hypothesis. Yet the agreement of the rate distributions, as well as other single cell spike statistics, indicate that these lingering correlations must in some sense be small, a sense that is yet to be precisely defined. But the essential point is that the effective smallness of these correlations is just as pertinent to experiment as to simulations (see e.g. [19]). It is rather safe to suppose that quantities that are insensitive to the CCs would be so in experiment as they are in the model, and may be captured by theory in both. Otherwise some weighty arguments on the position of the demarcation between the two should be advanced.

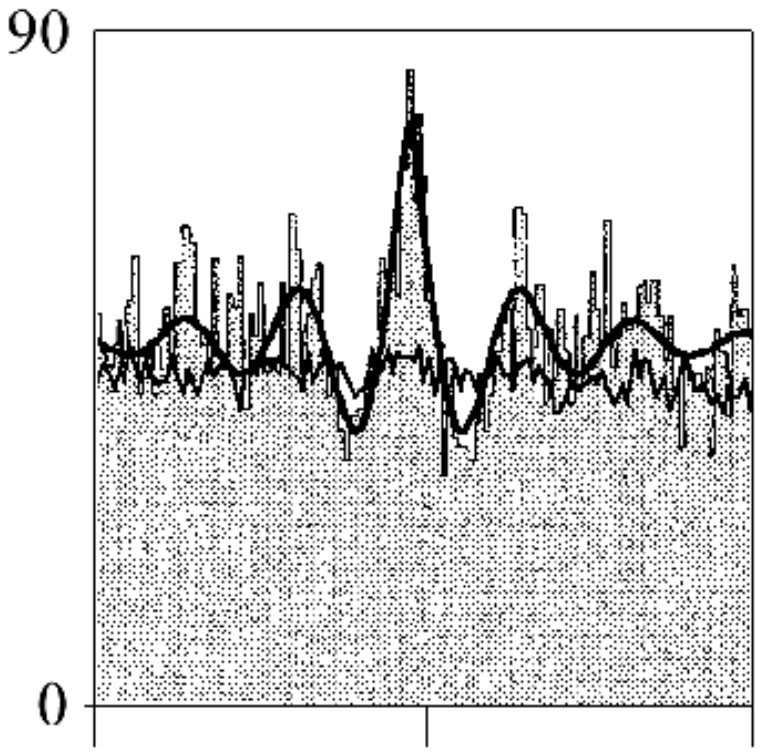

Figure 8: Experimental averaged cross-correlation from Figure 7 of Kreiter and Singer 1996[17]. Coincidence bin is 1ms. Horizontal axis (–64,+64ms); rates 50–70Hz. The oscillatory and noisy solid curves are theoretical superpositions

The secret of cross-correlations

In the last section cross-correlations appear as a thorn in the nice correspondence between the phenomenology of the simulated neural system and a very compact theory. But those same correlations highlight two other interesting issues. First, the simulated system makes predictions about what may be expected of the biological multi-unit recordings. Simulations indicate that CCs in the unstructured network are essentially independent of the efficacies, or even of the existence of a connection between the pair of neurons. See e.g. ref. [5]. On the other hand, CCs depend significantly on factors such as the rates of the two neurons and on the magnitude of transmission delays (or, alternatively, on synaptic integration time constants). Those are predictions, or tests of the biological adequacy of the hypotheses underlying the simulated system.

There is yet another aspect implied by the unsatisfactory relation between simulation and theory concerning CCs, which highlights both the critical and the constructive potential of the simulation. Since CCs are more and more frequently measured as multi-unit techniques spread, the question of the origin of such correlations, as a key to the information they reveal, is increasingly pressing. As was mentioned above, CCs do not relate in a direct way to the synapses connecting the two neurons, as was often assumed[20]. The simulation is a case in which all the constitutive information about the system is both known and under control. Our inability to decipher the origin of the CCs in that system leaves little credibility to any interpretation of the significance of CCs in cortex of a performing animal. This is the negative message. The converse of this coin is that given a system so exceedingly oversimplified (as any neuro-biologist would testify) it may allow a theoretical account of the CCs it produces. If such a theory emerges, it will constitute a credible hypothesis for CCs in cortical situations in-vivo. Existing theoretical analyses of CCs, while accounting quite well for their dependence on various parameters, fail to capture the magnitude of the central peak (e.g. ref. [21]).

Learning and delay activity – working memory

On the substrate of a credible spontaneous activity described above one may attempt some version of Hebbian learning. Here the need for a guiding system is even more urgent, since so much less is known about learning (synaptic dynamics) in-vivo than about neural spike activity. We have considered learning in the simple network of integrate-and-fire neurons presented above (refs. [15, 5]). It consists of the presentation of series of stimuli, as excitatory currents exciting subsets of the excitatory neurons of the network. Against a background of random synapses sustaining spontaneous activity, synapses are modified: they are strengthened if they connect two excited neurons and weakened if they connect an active to a passive neuron, maintaining the total synaptic efficacy in the network constant. This is a particular scenario of learning, for which theory has very definite predictions: first spontaneous activity should become inhomogeneous; then when potentiation has become strong enough, selective delay activity distributions, as in delay response experiments, can auto-sustain themselves and become correlated[22, 4].

Learning, and the persistence of delay activity distributions, may also be related to the lingering issue of the presence of a large peak near zero frequency in the rate distribution (Figure 7), absent in the theory and in the recordings from the simulation in pure spontaneous

activity states. Such enhancement in the number of quiescent cells may be due to attention that accompany discrimination experiments. It is a challange that should and can enrich the experimental-simulation-theoretical methodology.

Outlook

Simulating in detail large scale neural systems does not solve any problem in cognitive, or computational, neurophysiology. It does however allow the placing of explanatory proposals in a well defined and controlled context that is not constructed to express a specific function. On the other hand, what is expressed in a simulation, becomes a good candidate for a concise, predictive theoretical analysis. One may proceed to produce a long list of effects of interest in experiment. In fact, most hypothesized scenarios about cortical dynamics can be tested against such simulations, both as a feasibility test and as a guide to experiment. The list would be as long as the imagination of experimental neurobiology. For example, one may test the response of the network to external stimuli, against the background of spontaneous activity. Stimuli may be simple injections of current into separate neurons, or may include effects of receptive fields and tuning curves. Such "experiments" may consist of measuring response times of the system (as in [23]); of CCs and their relation to peristimulus phenomena (as in ref. [17]). The model may be extended to test the effects of long range inhibition, as indicated in[18]. The simulated neural dynamics may suggest a role for non-stationary action of potentiated synapses [24]. It may propose a clear role for synaptic conductance time constants (not included in the simulations described here). Simulation techniques are now on the feasibility verge of joint neural and synaptic dynamics (learning) in large-scale networks, which may allow a more organic connection of the modeled cortical module to the external world. They may and should be extended to include interactions between modules, thus introducing the computational dynamics of brain-like systems. All these phenomena would be expressed and measured on a well understood substrate and proposed accounts can be put to a test, whereas in the living system they are often left as untestable metaphors. Moreover, given the simplicity of the system they may find a theoretical, predictive account.

Acknowledgements

I am deeply grateful to Dr Paul Ginsparg for a a critical reading of an earlier version of the manuscript which left deep traces. I thank Dr Volodya Yakovlev for help with the experimental data and Dr Nicolas Brunel for discussions and simulation data and Dr Sergei Gepstein, Misha Dvorkin and Stefano Fusi for help with the graphics. This work was supported in part by a Human Mobility grant of the EEC and by grant 509/96 from the Israel Science Foundation.

References

[1] Rall W., Perspective on Neuron Model Complexity, *The Handbook of Brain Theory and Neural networks* in Arbib MA, ed. MIT press 1995 and references therein.

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

[3] DeSchutter, E., and Bower, J.M. 1994 An Active Membrane Model of the Cerebellar Purkinje Cell: I simulation of current clamps in slice, *J. Neurophysiol.* **71** 375

[4] D.J. Amit, N. Brunel and M.V. Tsodyks 1994 Correlations of Hebbian reverberations: theory vs experiment, *Jour. Neurosci.* **14** 6435

[5] D.J. Amit, N. Brunel 1997 Dynamics of a recurrent network of spiking neurons before and following learning - *NETWORK Computation in Neural System* in press Oct 1997. (See e.g. www.fiz.huji.ac.il/staff/acc/faculty/damita)

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

[7] Jefferys, J.G.R., Whittington, M.A. and Traub, R.D. 1996 Neuronal networks for induced 40Hz rhythms, TINS, **19** 202-208

[8] D. Badoni, S. Bertazzoni, S. Buglioni, G. Salina D. J. Amit and S. Fusi (1995) Electronic implementation of an analog attractor neural network with stochastic learning, *NETWORK* **6** 125

[9] F. van der Velde 1996 Integrating connectionism and symbol manipulation: the importance of implementation, *Behavioral and Brain Sciences* to be published

[10] Miyashita Y and Chang HS 1988 Neuronal correlate of pictorial short-term memory in the primate temporal cortex, *Nature*, **331** 68

[11] Ricciardi LM 1977 *Diffusion processes and Related topics on biology* (Springer-Verlag, Berlin)

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

[13] Amit DJ, Fusi S and Yakovlev V 1997 A paradigmatic working memory (attractor) cell in it cortex *Neural Computation*, **9** 1101

[14] Yakovlev V, Fusi S and Dvorkin M 1997 unpublished.

[15] 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** 237

[16] van Vreeswijk C and Sompolinsky H 1996 Chaos in neuronal networks with balanced excitatory and inhibitory activity, *Science* **274** 5293

[17] 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

[18] Kisvarday ZF, Kim DS, Eysel UT and Bonhoeffer T 1994, Relationeship between lateral inhibitory connections and topography of the orientation map in cat visual cortex *Eur. J. Neurosci.,* **6** 1619

[19] Amit DJ 1997, Is synchronization necessary and is it sufficient? Peer commentary on Phillips WA and Singer W 1997 In search of common foundation for cognitive science, *Behavioural and Brain Sciences* In press.

[20] 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

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

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

[23] Thorpe S Fize D and Marlot C 1996, Speed of processing in the human visual system, *Nature* **381** 520

[24] Markram H and Tsodyks MV 1996 Redistribution of synaptic efficacy between neocortical neurons *Nature* **382** 807