

Transient information flow in a network of excitatory and inhibitory model neurons: Role of noise and signal autocorrelation

Julien Mayor *, Wulfram Gerstner

*School of Computer and Communication Sciences and Brain-Mind Institute, Ecole Polytechnique Fédérale de Lausanne,
CH-1005 Lausanne EPFL, Switzerland*

Abstract

We investigate the performance of sparsely-connected networks of integrate-and-fire neurons for ultra-short term information processing. We exploit the fact that the population activity of networks with balanced excitation and inhibition can switch from an oscillatory firing regime to a state of asynchronous irregular firing or quiescence depending on the rate of external background spikes.

We find that in terms of information buffering the network performs best for a moderate, non-zero, amount of noise. Analogous to the phenomenon of stochastic resonance the performance decreases for higher and lower noise levels. The optimal amount of noise corresponds to the transition zone between a quiescent state and a regime of stochastic dynamics. This provides a potential explanation of the role of non-oscillatory population activity in a simplified model of cortical micro-circuits.

© 2005 Elsevier Ltd. All rights reserved.

Keywords: Recurrent integrate-and-fire neuron networks; Sparse connectivity; Population dynamics; Information processing

1. Introduction

The brain processes information in a constantly varying world. The visual scenes one sees in everyday life are extremely rich and change rapidly in time. In addition, the human eye performs more than three saccades every second causing sudden changes in the visual input [53]. Complex sounds, such as speech or music vary continuously in time and frequency. Moreover, the human brain has to deal with simultaneous sound sources from different origins, that are superimposed and mixed together. From such a perspective, it is clear that cortical areas are confronted with time-varying sensory inputs rather than stationary stimuli. While attractor neural networks are considered suitable as models of working memory [67,22,1], they are less useful to explain ultra-short memory buffer properties in the range of 10–100 ms of signals that vary continuously

in space or time. Recently, models of continuous information processing in recurrent neural networks have been proposed under the names of liquid state machine (LSM, [45]) and echo state networks [31,32], both related to the timing network of Buonomano and Merzenich [9]. Those models, that perform computation using transient activity (as opposed to convergence to a stable state), are sometimes referred to as models of computation with dynamic states. The idea underlying those models is that the instantaneous state of the network provides a rich reservoir of non-linear spatio-temporal transformations of the inputs.

In the framework of the liquid state machines, learning only acts on the readout structures, the network itself remaining fixed. Several recent studies focus either on the biological realism of such model networks or on possible implementations of this approach in machine learning or robotics. The role played by dynamical synapses was studied from the perspective of efficient temporal processing in [50,51]. Principles of liquid state machines were applied to the analysis of a variety of time series [49,52] for computer vision [44], movement generation and control [33,39], and

* Corresponding author.

E-mail address: julien.mayor@epfl.ch (J. Mayor).

URL: <http://diwww.epfl.ch/~jmayor> (J. Mayor).

prediction of chaotic time series [32]. The idea of performing (simple) computations based on perturbation of a real liquid (water) was investigated in [21].

Cortical micro-circuits are extremely complex recurrent networks of neurons. A given neuron is functionally connected to only a relatively small fraction of the other neurons. The connection strength between neurons is not fixed, but is mediated by synapses that have their own dynamics [3,56,15,19,28,60]. Maass and colleagues have chosen to simulate networks with a detailed set of biologically-inspired parameters. They implemented distance-dependent connectivity and different refractory periods and thresholds for the inhibitory and excitatory pools. They also investigated the role of dynamic synapses and introduced stochasticity in the values of some parameters [45].

On the other end of the spectrum of neural network models are idealised networks that replace real cortical connectivity by sparse random connections and reduce neuronal diversity to two neuron types, i.e. excitatory and inhibitory. Even from these simplified models, a rich dynamics can emerge such as fast and slow oscillations, synchrony and even chaotic behaviour [5,7,29,62,66]. These simplified models have the advantage that theoretical methods for investigating the dynamics are available via an analytical approach, so that the parameter space can be explored systematically. Based on mathematical analysis, the network behaviour can be classified in a small number of types, such as fast oscillations, slow oscillations, or spontaneous asynchronous activity [5,20], a rigorous classification which would be difficult to obtain by pure simulation-based studies.

A remarkable result of such studies concerns asynchronous activity [5,7]: in the phase diagram of the activity of sparsely-connected network of inhibitory and excitatory neurons, there exists a stable phase of asynchronous irregular firing in which the overall activity is stationary (its statistical properties such as mean and variance are time-independent) but the activity of individual neurons is highly irregular. From the point of view of information processing in a cortical-like network, such a regime of firing provides very interesting properties. In particular, we can imagine that information about past input can be buffered in the perturbations of the stable state of asynchronous activity.

The presence of such a memory trace of past stimulations is hypothesised for example in models of conditioning (such as reinforcement learning [64]) under the name of eligibility trace. However the underlying mechanism for such a trace is not known, even though many models have blossomed in the past decade. In most models (for example [63]), the clock is replaced by a series of spectra without any strong biological relevance. Bullock et al. [4] have a model of a clock that is dependent on the Ca^{2+} activity and Contreras-Vidal [13] hypothesised an activation of different subsets of striosomes. However these mechanisms are supposed to cover many orders of magnitude in timing (from a fraction of a second up to a few seconds), and it is

unlikely that such a broad spectrum can be covered by a single ionic mechanism. Another way of keeping track of time is to have a recurrent neural network, that will store dynamically the timing information. Such a mechanism is likely to exist in the olivo-cerebellar system [34].

In this paper we make use of the mathematical analysis of sparsely-connected networks [5] in order to relate macroscopic states of the network with evaluations of information buffering, thus establishing a link between network theory and the liquid state machines [45]. Specifically, we inject a time-dependent input current into a network of excitatory and inhibitory neurons with sparse connectivity. We want to know for how long information about an input signal is kept in the network. In order to get a quantitative answer, we attempt to reconstruct the input signal at time $t - T$ using the instantaneous membrane potentials of all neurons at time t . If the error of signal reconstruction is small, we say that the network is capable of “buffering” information for a time T .

2. Model and methods

2.1. Integrate-and-fire neuron

The system we study is a sparsely connected network of leaky integrate-and-fire (IF) neurons. Our network is composed of $n = 200$ IF neurons, 80% of which are excitatory and 20% inhibitory. The equation for the membrane potential of such neurons can be written as

$$\tau_m \dot{u}_i(t) = -u_i(t) + RI_i(t) \quad (1)$$

where I_i is the total input, $\tau_m = RC_m$ the effective membrane time constant, R the effective input resistance and C_m the membrane capacity. The total input of neuron i can be decomposed into contributions of the presynaptic spikes from other neurons j within the micro-circuit under consideration, an external drive I^{ext} and a noise term I^{noise} that models background spikes (described in detail in Section 2.2) from other brain areas that are not described explicitly

$$RI_i(t) = \tau_m \sum_{j \in M_i} \omega_j \sum_k \delta(t - t_j^k - D) + RI_i^{\text{ext}}(t) + RI_i^{\text{noise}}(t) \quad (2)$$

where M_i is the ensemble of presynaptic neurons, t_j^k the time neuron j fires its k th spike and D is a transmission delay. Synaptic weights are chosen from two values, either excitatory, $\omega_j = \omega_E$ or inhibitory, $\omega_j = -\omega_I$. Thus from Eqs. (1) and (2), the EPSP's resulting from spike arrivals at an excitatory synapse are of amplitude ω_E and decay with the membrane time constant τ_m . Analogously, the IPSP's resulting from spike arrivals at an inhibitory synapse are of amplitude $-\omega_I$ and decay with the membrane time constant τ_m . The total depolarisation of neuron i is the sum of all EPSP's and IPSP's that have arrived since the last spike of postsynaptic neuron i .

Table 1
(Left) numerical values of the model neurons (leaky integrate-and-fire neuron) used for the simulations

Neuron parameters			Network parameters		
Membrane time constant	τ_m	20 ms	Number of neurons	n	200
Threshold	θ	5 mV	Connectivity ratio	ϵ	0.2
Absolute refractory period	τ_{rp}	2 ms	Excitatory synaptic efficacy	ω_E	1 mV
Transmission delay	D	1 ms	Inhibitory synaptic efficacy	ω_I	5 mV
Reset membrane potential	u_{reset}	0 mV	Number of external excitatory neurons	N_E	100
Effective input resistance	R	10 M Ω			
Membrane capacity	C_m	2 nF	Spiking freq. of the external excitatory neurons (model A)	ν_{exc}	10^{-5} –100 Hz

Voltage is measured with respect to the resting potential. (Right) Numerical parameters of the simulated neural network. The connection probability of one neuron to another corresponds to the connectivity ratio ϵ .

When the depolarisation $u_i(t)$ of neuron i reaches a threshold θ , a spike is emitted and the membrane potential is reset to a potential u_{reset} after an absolute refractory period τ_{rp} during which the neuron is insensitive to any stimulation. The numerical values are shown in Table 1.

2.2. Stochastic background input

We consider two noise models. The first model assumes stochastic arrival of *excitatory* spikes from other areas. Since a complete analytical description of the impact of this noise model on the dynamics of the network has been done [5], we can connect our results to the known macroscopic dynamics of randomly connected networks.

In a second noise model, we include both excitatory and inhibitory spikes from outer regions. The combination of inhibition with excitation generalises the first model and permits change in the fluctuations of the drive without changing its mean. It also allows us to draw an interesting link to a physical phenomena known under the name of *stochastic resonance*.

2.2.1. Model A: stochastic spike arrival at excitatory synapses

In order to simulate the intense synaptic bombardment of neurons observed in vivo (for a review see [14]), we first consider purely excitatory stochastic spike arrivals coming from N_E external excitatory neurons. We assume that each spike changes instantaneously the membrane potential of the post-synaptic neurons (satisfying Eq. (1)). Stochastic spike arrivals can be described by the sum of N_E homogeneous Poisson processes of rate ν_{exc} . In such a description the probability of having a spike in a time window is stationary in time. There is no correlation between spikes, i.e. the occurrence of a spike at a given time does not influence the future. If we make the approximation that a neuron receives a large number of presynaptic contributions per unit time, each generating a change ω_E in the membrane potential that is relatively small compared to the firing threshold θ , the noise term can be written as a constant drive plus a fluctuating part:¹

$$RI_i^{noise}(t) = \mu_E + \sigma_E \sqrt{\tau_m} \eta_i(t) \quad (3)$$

where $\mu_E = N_E \omega_E \nu_{exc} \tau_m$ and $\sigma_E = \omega_E \sqrt{N_E \nu_{exc} \tau_m}$. By η , we denote standard Gaussian white noise of zero mean and unit variance, i.e.,

$$\begin{aligned} \langle \eta_i(t) \rangle &= 0 \\ \langle \eta_i(t) \eta_i(t') \rangle &= \delta(t - t') \end{aligned}$$

This type of excitation correspond to the analysis of asynchronous activity in sparsely connected network of IF neurons made in [5]. The results from our simulations can thus be directly interpreted with the phase diagram of the activity, Fig. 2 in Ref. [5].

We will focus on the influence of the external spiking activity to the performance of the network in the range 10^{-5} Hz $<$ ν_{exc} $<$ 100 Hz.

2.2.2. Model B: excitatory and inhibitory stochastic spike arrivals

In the second noise model we assume stochastic spike arrival from both excitatory and inhibitory neurons. Let us consider that we have an external population N_E of excitatory neurons firing at a rate ν_{exc} , as in model A. We now add a further input population (referred thereafter as the “balanced population”) made out of two groups of neurons, one of them excitatory and the other one inhibitory. In order to have a balance between excitation and inhibition, we choose to have five times as many excitatory neurons as inhibitory neurons, but with inhibitory synapses that are five times as strong, i.e., $N_E^{pop} = \frac{\omega_I}{\omega_E} N_I^{pop}$, with $\frac{\omega_I}{\omega_E} = 5$. This approximate balance between excitation and inhibition is thought to take place at a functional level in cortical areas [41,59]. The choice we made to have an exact balance in this additional population of neurons is for the sake of simplicity only. Its mean contribution is therefore zero and a change in the firing rate of this population only affects the variance of the drive. In addition, modulation of the variance part of the drive will allow us to focus on pure noise effects for a fixed mean drive.

With such a model we can write the noise term, similar to the previous model A

$$RI_i^{noise}(t) = \mu_E + \sigma_{pop} \sqrt{\tau_m} \eta_i(t) \quad (4)$$

where we have defined $\mu_E = N_E \omega_E \nu_{exc} \tau_m$ and $\sigma_{pop} = \sqrt{\tau_m [N_E \omega_E^2 \nu_{exc} + \nu_{pop} (N_E^{pop} \omega_E^2 + N_I^{pop} \omega_I^2)]}$.

The effective mean drive corresponds to the mean excitation from the purely excitatory population only since the

¹ This substitution is known as a diffusion approximation, see e.g. [5].

mean contribution of the balanced population is zero. The variance can now be varied independently of the effective mean input, by changing the discharge rate ν_{pop} of the balanced population.

The noise terms in both models (A) and (B) can be expressed as:

$$RI_i^{\text{noise}}(t) = \mu_E + \sigma\sqrt{\tau_m}\eta_i(t)$$

where $\sigma = \sigma_E$ in noise model (A) and $\sigma = \sigma_{\text{pop}}$ in noise model (B). Since the drive $\mu_E = N_E\omega_E\nu_{\text{exc}}\tau_m > 0$ corresponds to the background activity of an excitatory population, the mean drive is always positive. For $N_E = 100$, $\omega_E = 1$ mV, $\nu_{\text{exc}} = 1$ Hz and $\tau_m = 20$ ms we get a mean depolarisation of $\mu_E = 2$ mV due to the noise input.

2.3. Network structure

Both excitatory and inhibitory neurons are modelled with a membrane time constant τ_m of 20 ms. The neurons are weakly (connection probability = 0.2) and randomly connected through simple static synapses. We carefully chose the synaptic strengths, $\omega_I = 5\omega_E$ so that for $\nu_{\text{exc}} = 2\text{--}5$ Hz and $N_E = 100$, the network is in a regime of asynchronous and irregular firing (see Fig. 3B centre), based on the phase diagram described in [5]. In such a regime, there is a stationary overall activity and highly irregular spike firing of individual cells. Specifically, we take $\tau_{\text{rp}} = 2$ ms, $D = 1$ ms, $\theta = 5$ mV, $u_{\text{reset}} = 0$ mV, $\omega_E = 1$ mV and $\omega_I = 5$ mV (see Table 1), but other combinations of parameters give qualitatively similar results. An extensive scanning of the parameter space has been carried out.² This robustness originates from the fact that the phase diagram of the type of networks we study in this article, can be qualitatively drawn as a function of two parameters only; the external drive (mean and variance) and the effective ratio between excitation and inhibition. The presence of an asynchronous phase of neural activity (and therefore the essence of our results) is guaranteed for a broad range of neuronal parameters [5].

With a stronger input, the system reaches a phase of synchronous firing with fast oscillations [5]. For weak drive, activity will tend to zero (quiescent state) since the network itself is not capable of sustaining intrinsic activity.

As we will see later, the absence of synchrony at the working point of the network is important in order to let information about a past stimulation flow for a long time in the network.

2.4. Evaluating information buffer properties

We perform simulations of the network with some time-dependent inputs $I_i^{\text{ext}}(t)$. The question we want to ask is:

² For example, the precise values of the synaptic strengths affect the exact location of the optimal performance when larger than 0.1 mV but not the overall trend of the results.

how much information is left about the external inputs after some time T ? Given some continuous input $I_i^{\text{ext}}(s)$ we expect that the instantaneous state of the network (given by its membrane potentials $\{u_j(t) | 1 \leq j \leq n\}$) holds information about $I_i^{\text{ext}}(s)$ for $s < t$. If a reconstruction of the input signal $I_i^{\text{ext}}(t - T)$ can be achieved by looking at the membrane potentials at time t , we say that the network can “buffer” information for a time T .

We assess the information buffering capacity of the network with a procedure analogous to [45], [31]. We inject simultaneously N_s independent input signals to N_s disjoint groups of randomly chosen neurons (see Fig. 1A), every neuron receiving exactly one input; e.g. for $N_s = 4$ input signals we have four groups of 50 neurons in our network of 200 neurons. N_s readout structures have access to the membrane potentials of all neurons in the network at time t and use this information to estimate the input signal that was present at time $t - T$ (see Fig. 1B).

We consider two different methods

- Linear read-out.

The outputs of the readout structures are simple linear combinations of all the membrane potentials i.e.

$$\text{Output}(t) = \sum_k \alpha_k u_k(t) + \alpha_0 = \vec{\alpha}^T \cdot \vec{u} + \alpha_0$$

where $\vec{u} = (u_1, u_2, \dots, u_n)^T$ is the vector of membrane potentials and α_0 a bias term.

- Non-linear kernel-based read-out.

In order to know whether more complex read-out structures can extract significantly more information than the linear read-out, we used support vector machines (SVM's) [10], a powerful method from the field of statistical learning [65]. It belongs to the family of kernel-based methods [58]. In those techniques, the input is sent to a high-dimensional space in which linear fits are expected to work, i.e.

$$\text{Output}(t) = \sum_l \alpha_l \mathcal{K}(\vec{u}(t), \vec{u}(t_l)) + \alpha_0$$

The sum runs over several samples of instantaneous network states $\vec{u}(t_l)$. The set of relevant samples is selected by the SVM algorithm. Different kernels $\mathcal{K}(\cdot)$ can be used in order to implicitly build a multi-dimensional space. In Section 6 we will compare the linear read-out to a support vector machine with a Gaussian kernel. We used SVM-torch [12] for the SVM method.

Only the parameters (the weights α_k) of the readout structures are tunable, while the network of IF neurons itself remains fixed. Since we want to assess how much we can learn about the input at time $t - T$ if we read out the membrane potentials at time t , we minimise the error:

$$E(\vec{\alpha}) = \frac{1}{2A} \sum_t [\text{Output}(t) - \text{Input}(t - T)]^2 \quad (5)$$

with $A = \sum_t (\text{Input}(t) - \bar{I})^2$ where \bar{I} denotes the mean drive. Minimisation is achieved by an optimisation proce-

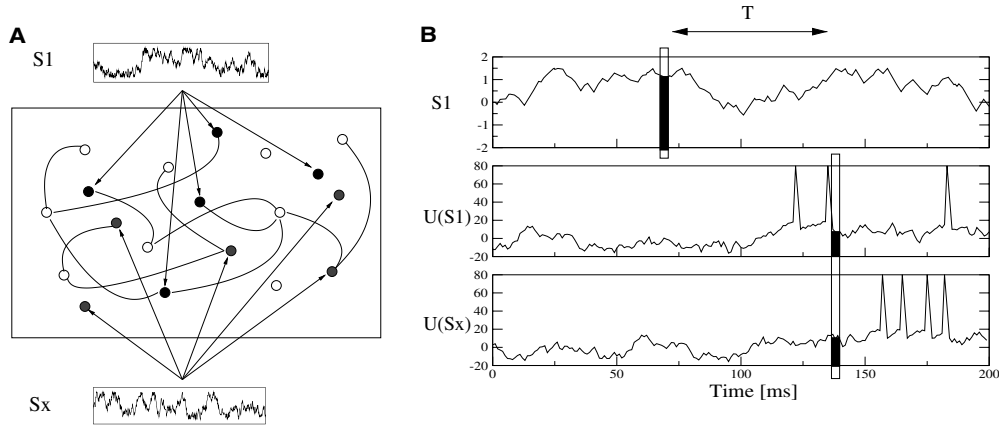


Fig. 1. (A) The different signals S_1, \dots, S_x are introduced to randomly chosen interconnected neurons. Every neuron receives exactly one input from the exterior. Only two input signals S_1 and S_x are shown in the figure. (B) Based on the momentary state of all membrane potentials (those receiving directly the signal S_1 and those receiving any other signal S_x), a readout structure is trained to guess the amplitude of its corresponding input a time T before (referred as delay thereafter).

procedure (optimal regression for the linear read-out) on the training set. Later in the results section, we will refer to the error defined in Eq. (5) as the *signal reconstruction error*. The input $(t - T)$ plays the role of a target value for optimisation. For the linear readout there is a unique optimal set of weights $\{\alpha_k\}$ for every delay T .³ After a training period, the weights of the readout structures are frozen and N_s new input signals are introduced in the network. Outputs of the readout structures are then compared with their corresponding targets. For all the simulations, the trajectories used for training correspond to 50,000 time steps of simulation. Performance is measured on an independent test set of 50,000 time steps. For the sake of clarity, all figures in the result section are obtained for a single input $N_s = 1$. In the discussion section, we vary the number N_s of inputs.

2.5. Autocorrelated inputs

As mentioned before, all neurons in the network receive one of N_s input signals $I_i^{\text{ext}}(t)$. We build an input such that its autocorrelation profile depends on a single parameter and decays to zero beyond a given value. In order to do so, the total simulation time is broken into segments of duration T_{max} . During each segment of length T_{max} the input is kept constant. At the transition to the next segment, a new value of $I_i^{\text{ext}}(t)$ is drawn from a uniform distribution centred at zero, between -20 pA and $+20$ pA (-7.5 pA and $+7.5$ pA as an additional example shown in Fig. 2B).

³ The application of the projection theorem [43] says that it is true if the number n of free parameters (the number of neurons, $n = 200$ in our case) is smaller than the number of independent examples. We took 50,000 examples, but these are slightly correlated in time. A realistic measure is to divide the number of examples by the effective autocorrelation of the input (3.3 ms up to 80 ms, see Section 2.5). Even in the worst case, a lower bound for the number of linearly independent examples is $\frac{50,000}{80} = 625 > n = 200$. The lengths of the simulations guarantees that there is a unique optimal set of weights.

This procedure results in a triangular autocorrelation profile $A(s)$ that is zero for $t > T_{\text{max}}$, i.e.

$$A(s) = \begin{cases} A_0(T_{\text{max}} - |s|) & \text{for } |s| < T_{\text{max}}, \\ 0 & \text{for } |s| \geq T_{\text{max}} \end{cases}$$

Thus the signal at time t provides no information about the signal at $t - T$ for $T > T_{\text{max}}$. With the triangular autocorrelation, we can compute the effective autocorrelation τ_{in} (that can be seen as the time constant) of such an input

$$\tau_{\text{in}} = \frac{\int_0^\infty |s|A(s) ds}{\int_0^\infty A(s) ds}$$

which yields $\tau_{\text{in}} = \frac{T_{\text{max}}}{3}$. In our simulations we explored values of τ_{in} in the range of $3.3 < \tau_{\text{in}} < 80$ ms.

2.6. Spectral analysis

For the interpretation of some of our results, it is useful to consider the power spectrum of the network activity. The spectrum is obtained by a fast Fourier transform on the temporal evolution of the spiking activity during 2^{16} time steps. The spectrum is then averaged over 1000 frequency steps in order to smooth the curves. Such spectral curves give us the power at different frequencies. A single peak at a given frequency would indicate for example oscillations in the network at that precise frequency. As an artifact of the averaging procedure sharp resonance peaks would be slightly broadened. We have checked that for the simulations shown in this paper, averaging does not cause artificial broadening.

3. Results

We stimulate a network of sparsely connected excitatory and inhibitory integrate-and-fire neurons by a continuous time-dependent input current $I^{\text{ext}}(t)$. We will evaluate for how long information about an input signal is kept in the




Fig. 2. (A) Power spectra of the population firing activity. (B) Error as a function of Poisson noise. Notice the non-monotonic dependence on the amount of noise. As explained in [5], the asynchronous firing phase can be found for an intermediate drive (rather flat power spectrum). In a lower noise regime the neurons stand farther from threshold, the network remains silent in absence of a signal. Hence a spike-based readout would only become effective for an external drive larger than $v_{\text{exc}} \simeq 0.8$ Hz. The zone of optimality is then located at the upper limit of the quiescent regime. In the high excitation limit, the mean drive dominates the effect of the “pure noise” component and oscillations appear.

network. To this end, we try to reconstruct the input signal at time $t - T$ by looking at the instantaneous membrane potentials of all neurons at time t . If the reconstruction is possible, we say that the network is capable of “buffering” information for a time T . The “buffering” performance as a function of T is the central quantity of interest. Information buffering is possible if the error of signal reconstruction is small.

In the first subsection we investigate the role of noise on the performance measured in terms of the signal reconstruction error. We show that a moderate amount of noise systematically improves the performance. The role of asynchronous firing is emphasised.

In the second subsection we focus on the temporal characteristics of the input signal. We derive an empirical relation between the temporal memory of the network and the effective autocorrelation of the input.

Finally we show that the rich dynamics of the network in its asynchronous state provides a good representation of its past stimulation. Therefore simple linear readouts are powerful enough to extract information and perform nearly as well as advanced kernel-based methods [58,10].

3.1. On the importance of noise

3.1.1. Stochastic arrival of excitatory background spikes

Networks of integrate-and-fire neurons are known to have complex dynamics [5,7]. In particular, our network of 200 excitatory and inhibitory neurons switches from highly synchronised neuronal activity to a very irregular firing regime, depending on the external input; Fig. 3A and B, see also [5]. In the asynchronous phase, individual cells have a coefficient of variation C_V^4 close to or larger than 1, reflecting this stochasticity (for the analytical com-

putation of C_V , see [5]). The asynchronous firing phase can be reached with an external noise term corresponding to stochastic spike arrival at excitatory synapses (noise model A, see methods). The phase diagram of such a network is fully described in [5].

In order to characterise the network activity with this first model of noise (Eq. (4) see methods), we perform a preliminary series of simulations for different stochastic spike arrival rates v_{exc} at excitatory synapses in absence of a deterministic external stimulus. The power spectrum of the activity illustrates the regime of firing for the different drives, in particular the regimes of fast oscillations (above $v_{\text{exc}} \simeq 10$ Hz, e.g. for $v_{\text{exc}} = 52.5$ Hz, see Fig. 2A) and irregular asynchronous firing ($v_{\text{exc}} \simeq 1 - 5$ Hz). In the asynchronous firing phase, the power spectrum has no significant resonance peak (Fig. 2A, $v_{\text{exc}} = 1.6$ Hz), the individual spike trains are irregular (Fig. 3A centre) and the overall activity has no oscillations (Fig. 3B centre) in absence of input signal.

In a second series of simulation, we apply a time-dependent stimulus current I_i^{ext} . Neuronal activity reflects the temporal structure of the input (see Fig. 3C). In order to assess whether we can estimate the input at time $t - T$ from the set of membrane potentials at time t , we measure the signal reconstruction error (see methods).

In Fig. 2B, short-term memory buffer performance is measured in terms of the reconstruction error on the test set for a delay $T = 10$ ms by varying the rate v_{exc} at $N_E = 100$ excitatory synapses. Information buffer properties are good (i.e. minimal reconstruction errors) up to $v_{\text{exc}} \simeq 2$ Hz. The zone of optimality extends approximately up to the transition zone between quiescent and asynchronous irregular firing phase. In the quiescent regime (up to $v_{\text{exc}} \simeq 1$ Hz), there is no activity in absence of the input signal I^{ext} as seen in the left part of Fig. 3B.⁵ However the neurons are close to threshold and the input signal,

⁴ This is a measure of the width of the distribution of spike intervals, defined as the ratio of the variance and the mean squared. A Poisson distribution has a value of $C_V = 1$. A value of $C_V > 1$, implies that a given distribution is broader than a Poisson distribution with the same mean. If $C_V < 1$, then the spike train is more regular than that generated by a Poisson neuron of the same rate (see for example [26]).

⁵ A more realistic readout, that would have access to the spikes of the individual neurons solely, would therefore only work when the network is not silent. Thus the zone of optimality would be dramatically narrowed.

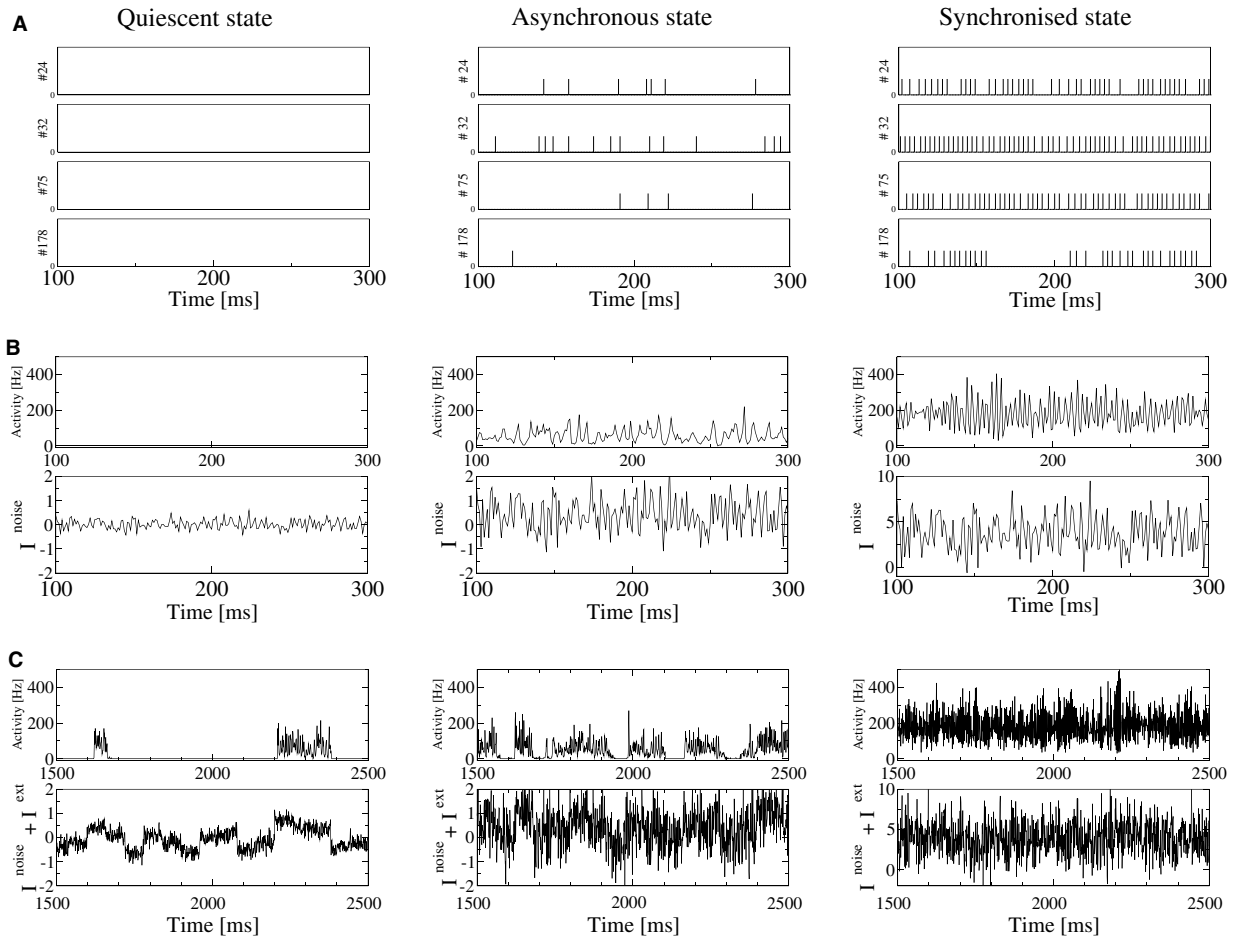


Fig. 3. (A) Representative spike trains in absence of the time dependent external signal $I^{ext}(t)$, for the quiescent state (left), for the asynchronous state (centre) and for the synchronised state (right). (B) Network activity and a representative noise signal for the quiescent state (left), for the asynchronous state (centre) and for the synchronised state (right). Note the different scale for the noise in the synchronised regime. (C) The external input $I^{ext}(t)$ has been switched on at time 1000 ms (not shown on the graph). We show the network activity and the sum of the external input signal $I^{ext}(t)$ and a representative noise signal for the quiescent state (left), for the asynchronous state (centre) and for the synchronised state (right).

that is switched on at time 1000 ms, make them fire when the input becomes sufficiently large, Fig. 3C left. For a weaker drive, neurons are farther from threshold and the input signal has to be stronger in order to elicit a response.

As the frequency ν_{exc} is increased beyond 1 Hz, the network reaches a phase of asynchronous firing. The population activity is now modulated by the input signal (Fig. 3C centre).

Increasing the excitatory background input does not only increase the noise, but also the mean drive. A resonance peak is already present in the high frequency domain for a moderate drive (Fig. 2A, $\nu_{exc} = 5$ Hz). The increased effective drive generates network activity at a high frequency (see the large peak at about 300 Hz in Fig. 2A, $\nu_{exc} = 52.5$ Hz). In such a regime, the individual cells fire at a large rate, Fig. 3A right, and the overall activity has fast oscillations in absence of input signal, Fig. 3B right. The excitation by the time-dependent external signal $I^{ext}(t)$ is buried in the noise of the stochastic excitatory background, and the overall activity is barely modulated by the signal,

Fig. 3C right. The performance hence decreases dramatically with an increase in the excitatory firing rate ν_{exc} .

We can thus deduce from the error plot in Fig. 2B, that the network operates at best close to the stochastic dynamics, rather than in a network dominated by oscillations.

Signature of chaos in neural systems is not a new concept. It was already shown in the past that networks of model neurons can give rise to a chaotic firing regime [29,62,66]. Experimental evidence indicates that cells exhibit in vivo very irregular spike trains, with a coefficient of variation C_V close to 1 (see e.g [61]), that may be the signature of a rich and very irregular underlying dynamics. From both an experimental and a theoretical perspective, the brain is thought to exhibit a chaotic behaviour [18,17,37].

In parallel, the fact that information processing by a dynamical system may gain from being close to a chaotic regime was postulated in the framework of cellular automata [38].

But only recently studies have focused on the fact that the brain may actually be operating at the edge of chaos [57,2]. Our findings suggest that networks of *spiking*

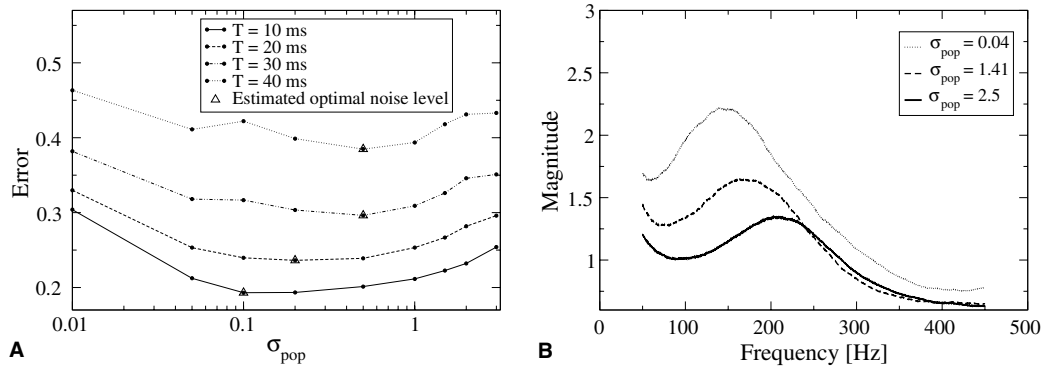


Fig. 4. (A) Errors as a function of the amplitude of noise in model B. Again, we can see that a moderate amount of noise level increases the performance of the network. In this case, this is purely a noise effect since the mean drive is constant. Thus it can also be interpreted as a stochastic resonance effect. Notice there is a different optimal noise level for the different delays T . A possible explanation is proposed in Section 3.1.2. (B) Corresponding power spectra of the activity. Since the constant drive is fixed, the noise just smoothes out the resonance peak present at a low noise level. Even though the peak in the power spectrum at high noise is slightly weakened, noise still has a corrupting effect on the input. The neurons are well desynchronised but the signal-to-noise ratio gets very low.

neurons can efficiently buffer information when close to a stochastic dynamics.

Spontaneous stochastic activity at low rates is also known to be a prerequisite of rapid signalling [25,6].

3.1.2. A relation to stochastic resonance: model B

We can also set the network in an state of a stationary network activity with irregular individual spike trains by driving it with a balanced excitatory and inhibitory spike input (see methods, noise model B). As detailed in Section 2.2.2, this drive is equivalent to a constant input (contributed by the purely excitatory population only) and additional white noise (the variance parts of both the purely excitatory population and the balanced population). The performance is evaluated for different white noise amplitude on the top of a constant depolarisation of $\mu_E = 0.6$ mV. Errors on the test set are computed for different delays T (see Fig. 4A). As previously seen for the noise model A, there is a non-monotonic dependence of the performance upon the noise level.

The signal reconstruction error is relatively high when the noise is low since the network is in a quiescent state. Addition of noise sets up the network in a state where the neurons are weakly coupled; each neuron act as an almost independent unit and its coupling to the other neurons enriches the network dynamics: we have reached a state in which the neurons are optimally desynchronised. By adding more noise, the signal-to-noise ratio decreases along with the ability to retrieve information. The resonance peak present in the low noise case is smoothed out as the noise level increases⁶ (see Fig. 4B). In contrast to Section 3.1.1, this non-monotonic dependence on the drive is now a pure noise effect, since the mean drive is kept constant.

⁶ In the case of a larger network, the resonance peak would be washed out. Since our simulations include a small number of neurons, finite size effects (such as the impact of fluctuations) become strong.

An additional interesting effect can be seen in Fig. 4A. For one delay T to another, the signal reconstruction error is at its lowest for a different noise level (referred as “optimal”). Surprisingly this optimal noise amplitude increases with the buffering time T . It means that if we want to set up the network so that information remains available for a buffering time of 40 ms, the noise has to be stronger than if we want to extract faithfully information after a buffering time of 10 ms. A possible explanation is that in the low noise limit, transients are reflected almost instantaneously in the activity profile [25,36]. A low noise level thus allows the readout to have a fine temporal resolution. In the high noise limit however, the network acts as a low pass filter [25,6]. Slow transients for high noise allow longer buffering times. Therefore the amplitude of noise has to be higher for long buffering times than for short ones.

Our results are related to the well-known phenomenon of stochastic resonance. Such an effect is seen when the response of a system to a drive depends non-monotonically on noise, with an optimum at a moderate, non-zero, noise level. There are many pointers in the literature to physical evidences of stochastic resonance in physical systems [69,11,23,24,47,71,70], and in models of neurons [8,42,55,54]. In living systems stochastic resonance has been reported in crayfish mechanoreceptors [16], the cricket cercal sensory system [40], neural slices [27], hippocampus [72], and the cortex [48]. More importantly the brain appears to actually make use of stochastic resonance at a functional level [35,30].

3.2. Impact of autocorrelated inputs

When the input is correlated in time, the readout can extract information from both the information buffered in the recurrent activity and from the input itself (filtered by the neurons). Thus if the autocorrelation of the input is short compared to the intrinsic temporal trace of the network, the role of the recurrency is crucial whereas for




Fig. 5. (A) Error plots as a function of time for different input statistics. An evident increase in performance is seen as the autocorrelation of the input gets longer. (B) Delays corresponding to an error of 0.5 as a function of the effective autocorrelation of the input. A linear fit shows us the clear dependence of the persistence of information in the network on the “self-memory” of the input.

very long autocorrelation, the output will predominantly read “online” the input, making use of the “self-memory”⁷ of the stimulation signal. We expect then to have an increase in performance as a function of the autocorrelation of the input. We can rule out an alternative explanation that a long autocorrelation has more time to impact the dynamics of the neural network by comparing the cross-correlation of the reconstructed output and the input with the signal autocorrelation for different autocorrelation profiles (for an illustrated example and additional arguments, see [46]). Thus the cross-correlation depends on the shape and on the duration τ_{in} of the signal autocorrelation.

We performed numerical simulations using different effective autocorrelation times τ_{in} for the input signals. The results are shown in Fig. 5A. We can estimate the effective information buffer by looking at the buffer time T at which the error rises to a value “arbitrary” chosen to be $E = 0.5$ (the similarity of the reconstructed output to the target is still visually evident).

As we can see in Fig. 5B, the simulation points can be fitted to first approximation by a straight line. In the limit of long autocorrelation $\tau_{in} \gg \tau_m$ we expect from simple mean field considerations, the buffer to have a temporal extension longer than $\frac{3}{2}\tau_{in}$, and to approach this line asymptotically from above.⁸ Indeed the simulation results

for large τ_{in} stay always above the lower bound represented by the dashed line in Fig. 5B.

For short autocorrelation profiles (of the order of the membrane time constant $\tau_m = 20$ ms or shorter), this approximation no longer holds, since the filtering effects due to the neuron’s membrane time constant comes into play. Rapid changes of the input on the scale of $\tau_{in} \ll \tau_m$ will be averaged out irremediably by the integration with the membrane time constant. Hence the network cannot buffer efficiently (i.e. with an error smaller than $E = 0.5$) input signals that have an effective autocorrelation τ_{in} smaller than about 5 ms (see Fig. 5B).

The performance can be improved for different autocorrelation profiles and for different network connectivities as proposed in [68]. They can also be increased by adding other time scales in the network, for example at the level of the synapses [45].

Analysis of cross-correlation and autocorrelation profiles (more precisely the analysis made in [46] of the location of the peak in the cross-correlation profile that is shifted in proportion of the autocorrelation length) show that the readout effectively makes an optimal trade-off between retrieving information from the autocorrelation of the signal and from the buffer provided by the network recurrency.

3.3. Simple linear readouts are efficient

As we have seen in the Section 3.1, injection of noise in the network helps in desynchronising the neurons. The resulting dynamics of the network is thus extremely rich. It has been shown that in the case of formal neural networks [2] and analogously for cellular automaton [38], the computational abilities of complex systems is at best when the system is close to a chaotic behaviour, as opposed to a dynamic state with short limit cycle. Another way of interpreting this complex behaviour is to notice that the network provides a broad set of different filters for the input. In other words, the one-dimensional input signal $I^{ext}(t)$ undergoes a “dimensional blow up”. This is a strong

⁷ On average the signal at time t holds information about itself at time $t - T$ for $T < T_{max}$.

⁸ In order to derive a lower bound we could use the averaged membrane potential for reconstruction instead all the membrane potentials of the 200 neurons in the network. Since the distribution of the membrane potentials is modulated by the external input I^{ext} , a reconstruction based on the average is possible. A trivial reconstruction can then be done assuming that the external input has not changed during the last time interval T . Assuming that the network can trivially buffer a signal that has not changed in the last time interval T , and that the network cannot reconstruct at all an input that has switched in this last time interval T , the effective buffer of the network for an input that changes every T_{max} is $T < \frac{1}{2}T_{max}$. Since the effective autocorrelation of our input is $\tau_{in} = \frac{1}{3}T_{max}$, the network has an effective buffer of $T = \frac{3}{2}\tau_{in}$.




Fig. 6. Comparison of performances between support vector machines (Gaussian kernel) and the simple linear readout on a test set. Optimisation has been performed on a separate training set of 50,000 time steps. A simple linear combination of the membrane potentials is extremely effective in extracting information and is comparable to that of a support vector machine.

correspondence to the working principle of some of the most efficient classification methods, named support vector machines (SVM) [10]. These methods, known under the generic name of kernel methods [58], send the input to a high dimensional space. In this multidimensional space, a simple linear separation (a hyper-plan) into two sub-spaces is able to perform the classification. An adapted version of those methods can handle regression problems. In Fig. 6 a comparison of the performances of both the simple linear readout and the SVM with a Gaussian kernel is represented (see methods).

Although the kernel method achieves slightly better results than the linear combination, the gain is of the order of a few milliseconds only. The explanation is that an additional dimensional blow up that is at the basis of kernel methods is not necessary since a sufficiently high dimensional representation is already given by the network itself. Hence a simple linear combination (corresponding to the hyper-plan in the kernel methods) is able to extract most of the information needed.

4. Discussion and conclusions

We investigated the buffering capacity of networks of excitatory and inhibitory integrate-and-fire neurons. The population firing activity of such a network with sparse connectivity can switch from an oscillatory firing state to an asynchronous irregular firing regime, depending on the rate of stochastic background spike arrival.

We found that the buffering capacity of the network is optimised in the presence of a moderate, non-zero, amount of noise. This non-trivial dependence upon noise can be related to the well-known phenomenon of stochastic resonance. The optimal noise level, or equivalently the optimal

discharge rate of an external population of neurons, corresponds approximately to the region of transition between a quiescent state and an asynchronous irregular firing regime. In this region, a very complex dynamics emerges from the network's activity. It provides a rich representation of the inputs. Simple adaptable linear readouts are thus able to extract the buffered information in an efficient way. Hence our results provide an interpretation for the role of non-oscillatory dynamics in a simplified model of cortical micro-circuits.

The buffering capacity of a such network can be generalised to the processing of a larger number of inputs. Simultaneous buffering of eight inputs is shown in Fig. 7 in comparison to a single input.

No significant degradation is seen for a moderate number of inputs. A loss in performance can be seen only when processing more than sixteen inputs (data not shown). This ability to handle multiple simultaneous inputs, along with the way information can be extracted from both the auto-correlation of the input signals and the buffer provided by the network recurrency, could provide an interesting way of combining different sources of information having different time scales. On somewhat more speculative grounds, we can think of the cortex as being in a state that is close to a chaotic behaviour, and that other neighbouring neural micro-circuits or even farther afferents may be tuning the *amount* of chaos (by changing their spiking rate for example) of a given cortical micro-circuit depending on the relevant task. One can also imagine to build artificial networks made out of such model neurons, in order to do both information processing and prediction of (even chaotic) time series. The present study would suggest to set up such a device in a state where its dynamics can easily switch to a more complex behaviour, depending on its drive.




Fig. 7. Error plots for simultaneous buffering of eight inputs. A comparison to the buffering of a single input (thick line) emphasises the fact that no significant degradation in performance is seen for a reasonable number of inputs. Our simulations show that a degradation is seen only from about sixteen simultaneous inputs.

References

- [1] D.J. Amit, N. Brunel, A model of spontaneous activity and local delay activity during delay periods in the cerebral cortex, *Cereb. Cortex* 7 (1997) 237–252.
- [2] N. Bertschinger, T. Natschläger, Real-time computation at the edge of chaos, *Neural Comput.* 16 (2004) 1413–1436.
- [3] V. Braitenberg, A. Schütz, *Anatomy of the Cortex*, Springer-Verlag, Berlin Heidelberg New York, 1991.
- [4] J. Brown, D. Bullock, S. Grossberg, How the basal ganglia use parallel excitatory and inhibitory learning pathways to selectively respond to unexpected reward cues, *J. Neurosci.* 19 (1999) 10502–10511.
- [5] N. Brunel, Dynamics of sparsely-connected networks of excitatory and inhibitory neurons, *J. Comput. Neurosci.* 8 (2000) 183–208.
- [6] N. Brunel, F. Chance, N. Fourcaud, L. Abbott, Effects of synaptic noise and filtering on the frequency response of spiking neurons, *Phys. Rev. Lett.* 86 (2001) 2186–2189.
- [7] N. Brunel, V. Hakim, Fast global oscillations in networks of integrate-and-fire neurons with low firing rates, *Neural Comput.* 11 (1999) 1621–1671.
- [8] A. Bulsara, E.W. Jacobs, T. Zhou, F. Moss, L. Kiss, Stochastic resonance in a single neuron model: theory and analog simulation, *J. Theor. Biol.* 152 (1991) 531–555.
- [9] D.V. Buonomano, M. Merzenich, Temporal information transformed into a spatial code by a neural network with realistic properties, *Science* 267 (1995) 1028–1030.
- [10] N. Christianini, J. Shawne-Taylor, *An Introduction to Support Vector Machines, and Other Kernel-Based Learning Methods*, Cambridge University Press, 2000.
- [11] J.J. Collins, C.C. Chow, A.C. Capela, T.T. Imhoff, Aperiodic stochastic resonance, *Phys. Rev. E* 54 (1996) 5575–5584.
- [12] R. Collobert, S. Bengio, Svmtorch: support vector machines for large-scale regression problems, *J. Mach. Learning Res.* 1 (2001) 143–160.
- [13] J. Contreras-Vidal, W. Schultz, A predictive reinforcement model of dopamine neurons for learning approach behavior, *J. Comput. Neurosci.* 6 (1999) 191–214.
- [14] A. Destexhe, M. Rudolph, D. Paré, The high-conductance state of neocortical neurons in vivo, *Nature Rev. Neurosci.* 4 (2003) 739.
- [15] R. Douglas, K. Martin, *Neocortex*, in: G. Shepherd (Ed.), *The Synaptic Organization of the Brain*, Oxford University Press, 1998, pp. 459–509.
- [16] J. Douglass, L. Wilkens, E. Pantazelou, F. Moss, Noise enhancement of information transfer in crayfish mechanoreceptors by stochastic resonance, *Nature* 365 (1993) 337–340.
- [17] P. Faure, H. Korn, A nonrandom dynamic component in the synaptic noise of a central neuron, *PNAS* 94 (1997) 6506–6511.
- [18] P. Faure, H. Korn, Is there chaos in the brain? I. Concepts of nonlinear dynamics and methods of investigation, *C. R. Acad. Sci. III* 324 (9) (2001) 773–793.
- [19] J.D. Felipe, Neocortical neuronal diversity: chemical heterogeneity revealed by colocalization studies of classic neurotransmitters, neuropeptides, calcium-binding proteins, and cell surface molecules, *Cereb. Cortex* 7 (1993) 476–486.
- [20] B. Ermentrout, Neural networks as spatio-temporal pattern-forming systems, *Rep. Prog. Phys.* 61 (1998) 353–430.
- [21] C. Fernando, S. Sojakka, Pattern recognition in a bucket, in: *Proc. of ECAL*, 2003.
- [22] S. Fusi, Hebbian spike-driven synaptic plasticity for learning patterns of mean firing rates, *Biol. Cybern.* 87 (2002) 459–470.
- [23] L. Gamaitoni, F. Marchesoni, S. Santucci, Stochastic resonance as a bona fide resonance, *Phys. Rev. Lett.* 74 (1995) 1052–1055.
- [24] L. Gamaitoni, P. Hänggi, P. Jung, F. Marchesoni, Stochastic resonance, *Rev. Mod. Phys.* 70 (1998) 223–287.
- [25] W. Gerstner, Population dynamics of spiking neurons: fast transients, asynchronous states and locking, *Neural Comput.* 12 (2000) 43–89.
- [26] W. Gerstner, W. Kistler, *Spiking Neuron Models. Single Neurons, Populations, Plasticity*, Cambridge University Press, 2002.
- [27] B.J. Gluckman, T.I. Netoff, E.J. Neel, W.L. Ditto, M.L. Spano, S. Schiff, Stochastic resonance in a neuronal network from mammalian brain, *Phys. Rev. Lett.* 77 (1996) 4098–4101.
- [28] A. Gupta, Y. Wang, H. Markram, Organizing principles for a diversity of gabaergic interneurons and synapses in the neocortex, *Science* 287 (2000) 273–278.
- [29] D. Hansel, H. Sompolinsky, Chaos and synchrony in a model of a hypercolumn in visual cortex, *J. Comput. Neurosci.* 3 (1996) 7–34.
- [30] I. Hidaka, D. Nozaki, Y. Yamamoto, Functional stochastic resonance in the human brain: Noise induced sensitization of baroreflex system, *Phys. Rev. Lett.* 85 (2000) 3740–3743.
- [31] H. Jaeger, Short term memory in echo state networks, *Tech. Rep. GMD Report 152*, German National Research Center for information technology (2002).
- [32] H. Jaeger, H. Haas, Harnessing nonlinearity: predicting chaotic systems and saving energy in wireless communication, *Science* 304 (2004) 78–80.
- [33] P. Joshi, W. Maass, Movement generation and control with generic neural microcircuits, in: *Proc. of the First International Workshop on Biologically Inspired Approaches to Advanced Information Technology 2004*, Lausanne, 2003.
- [34] W.M. Kistler, C.I. De Zeeuw, Dynamical working memory and timed responses: the role of reverberating loops in the olivo-cerebellar system, *Neural Comput.* 14 (2002) 2597–2626.
- [35] K. Kitajo, D. Nozaki, L.M. Ward, Y. Yamamoto, Behavioral stochastic resonance within the human brain, *Phys. Rev. Lett.* 90 (2003) 218103.
- [36] B.W. Knight, The relationship between the firing rate of a single neuron and the level of activity in a population of neurons, *J. Gen. Physiol.* 59 (1972) 767–778.
- [37] H. Korn, P. Faure, Is there chaos in the brain? II. Experimental evidence and related models, *C. R. Biol.* 326 (2003) 787–840.
- [38] C.G. Langton, Computation at the edge of chaos—phase-transitions and emergent computation, *Physica D* 42 (1990) 12–37.
- [39] R.A. Legenstein, H. Markram, W. Maass, Input prediction and autonomous movement analysis in recurrent circuits of spiking neurons, *Rev. Neurosci. (Special Issue on Neuroinform. Neural Artif. Comput.)* 14 (1–2) (2003) 5–19.
- [40] J. Levin, J. Miller, Broadband neural encoding in the cricket cercal sensory system enhanced by stochastic resonance, *Nature* 380 (1996) 165–168.
- [41] G.S. Liu, Local structural balance and functional interaction of excitatory and inhibitory synapses in hippocampal dendrites, *Nat. Neurosci.* 7 (4) (2004) 373–379.
- [42] A. Longtin, Stochastic resonance in neuron models, *J. Stat. Phys.* 70 (1993) 309–327.
- [43] D.G. Luenberger, *Optimization by Vector Space Methods*, Wiley, New York, 1997.
- [44] W. Maass, R.A. Legenstein, H. Markram, A new approach towards vision suggested by biologically realistic neural microcircuit models, in: H.H. Buelthoff, S.W. Lee, T.A. Poggio, C. Wallraven (Eds.), *Biologically Motivated Computer Vision. Proc. of the Second International Workshop, BMCV 2002*, Tübingen, Germany, November 22–24, 2002, *Lecture Notes in Computer Science*, 2525, Springer, Berlin, 2002, pp. 282–293.
- [45] W. Maass, T. Natschläger, H. Markram, Real-time computing without stable states: a new framework for neural computation based on perturbations, *Neural Comput.* 14 (11) (2002) 2531–2560.
- [46] J. Mayor, W. Gerstner, Online processing of multiple inputs in a sparsely-connected recurrent neural network, in: O. Kaynak, E. Alpaydin, E. Oja, L. Xu (Eds.), *Proc. Joint Int. Conf. ICANN/ICONIP 2003*, Springer-Verlag, Heidelberg, 2003, pp. 839–845.
- [47] B. McNamara, K. Wiesenfeld, Theory of stochastic resonance, *Phys. Rev. A* 39 (1989) 4854–4869.
- [48] T. Mori, S. Kai, Noise-induced entrainment and stochastic resonance in human brain waves, *Phys. Rev. Lett.* 88 (2002) 218101.

- [49] T. Natschläger, W. Maass, H. Markram, The liquid computer: a novel strategy for real-time computing on time series, Special Issue Found. Inf. Process. *TELEMATIK* 8 (1) (2002) 39–43.
- [50] T. Natschläger, W. Maass, E.D. Sontag, A. Zador, Processing of time series by neural circuits with biologically realistic synaptic dynamics, in: T.K. Leen, T.G. Dietterich, V. Tresp (Eds.), *Adv. Neural Inf. Process. Syst.* 2000 (NIPS '2000), 13, MIT Press, Cambridge, 2001, pp. 145–151.
- [51] T. Natschläger, W. Maass, A. Zador, Efficient temporal processing with biologically realistic dynamic synapses, *Network: Comput. Neural Syst.* 12 (2001) 75–87.
- [52] T. Natschläger, H. Markram, W. Maass, Computer models and analysis tools for neural microcircuits, in: R. Kötter (Ed.), *Neurosci. Databases A Practical Guide*, Kluwer Academic Publishers, Boston, 2003, pp. 123–138 (Chapter 9).
- [53] J.K. O'Regan, Solving the real mysteries of visual perception: the world as an outside memory, *Can. J. Psychol.* 46 (3) (1992) 461–488.
- [54] H.E. Plesser, W. Gerstner, Noise in integrate-and-fire models: from stochastic input to escape rates, *Neural Comput.* 12 (2000) 367–384.
- [55] H.E. Plesser, S. Tanaka, Stochastic resonance in a model neuron with reset, *Phys. Lett. A* 225 (1997) 228–234.
- [56] S. Ramón y Cajal, *Histologie du système nerveux de l'homme et des vertébrés*, A. Maloine, Paris, 1909.
- [57] D. Rowe, Dynamic neural activity as chaotic itinerancy or heteroclinic cycles? *Behav. Brain Sci.* 24 (2001) 827–828.
- [58] B. Schölkopf, A.J. Smola, *Learning with Kernels*, MIT Press, 2002.
- [59] Y. Shu, A. Hasentraub, A.A. McCormick, Turning on and off recurrent balanced cortical activity, *Nature* 423 (2003) 288–293.
- [60] G. Silberberg, A. Gupta, H. Markram, Stereotypy in neocortical microcircuits, *Trends Neurosci.* 25 (5) (2002) 227–230.
- [61] W. Softky, C. Koch, The highly irregular firing pattern of cortical cells is inconsistent with temporal integration of random epsps, *J. Neurosci.* 13 (1993) 334–350.
- [62] H. Sompolinsky, A. Crisanti, H. Sommers, Chaos in random neural networks, *Phys. Rev. Lett.* 61 (1988) 259–262.
- [63] R.E. Suri, W. Schultz, Learning of sequential movements with dopamine-like reinforcement signal in neural network model, *Exp. Brain Res.* 121 (1998) 350–354.
- [64] R. Sutton, A. Barto, *Reinforcement Learning*, MIT Press, Cambridge, 1998.
- [65] V.N. Vapnik, *The Nature of Statistical Learning Theory*, second ed., Springer Verlag, 1999.
- [66] C. van Vreeswijk, H. Sompolinsky, Chaos in neuronal networks with balanced excitatory and inhibitory activity, *Science* 274 (1996) 1724–1726.
- [67] X.-J. Wang, Synaptic reverberation underlying mnemonic persistent activity, *Trends Neurosci.* 24 (2001) 455.
- [68] O.L. White, D.D. Lee, H. Sompolinsky, Short-term memory in orthogonal neural networks, *Phys. Rev. Lett.* 92 (2004) 148102.
- [69] K. Wiesenfeld, F. Jaramillo, Minireview of stochastic resonance, *Chaos* 8 (1998) 539–548.
- [70] K. Wiesenfeld, F. Moss, Stochastic resonance and the benefits of noise: from ice ages to crayfish and squids, *Nature* 373 (1995) 33–36.
- [71] K. Wiesenfeld, D. Pierson, E. Pantazelou, F. Moss, Stochastic resonance on a circle, *Phys. Rev. Lett.* 72 (1994) 2125–2129.
- [72] M. Yoshida, H. Hayashi, K. Tateno, S. Ishizuka, Stochastic resonance in the hippocampal ca3-ca1 model: a possible memory recall mechanism, *Neural Network* 15 (10) (2002) 1171–1183.