

Modelling the formation of working memory with networks of integrate-and-fire neurons connected by plastic synapses

Paolo Del Giudice ^{a,*}, Stefano Fusi ^b, Maurizio Mattia ^a

^a *Physics Laboratory, Istituto Superiore di Sanità, v.le Regina Elena 299, 00161 Roma, Italy*

^b *Institute of Physiology, University of Bern, CH-3012 Bülhplatz 5, Switzerland*

Abstract

In this paper we review a series of works concerning models of spiking neurons interacting via spike-driven, plastic, Hebbian synapses, meant to implement stimulus driven, unsupervised formation of *working memory* (WM) states.

Starting from a summary of the experimental evidence emerging from *delayed matching to sample* (DMS) experiments, we briefly review the attractor picture proposed to underlie WM states. We then describe a general framework for a theoretical approach to learning with synapses subject to realistic constraints and outline some general requirements to be met by a mechanism of Hebbian synaptic structuring. We argue that a stochastic selection of the synapses to be updated allows for optimal memory storage, even if the number of stable synaptic states is reduced to the extreme (bistable synapses). A description follows of models of spike-driven synapses that implement the stochastic selection by exploiting the high irregularity in the pre- and post-synaptic activity.

Reasons are listed why *dynamic learning*, that is the process by which the synaptic structure develops under the only guidance of neural activities, driven in turn by stimuli, is hard to accomplish. We provide a ‘feasibility proof’ of dynamic formation of WM states, by showing how an initially unstructured network autonomously develops a synaptic structure supporting simultaneously stable spontaneous and WM states in this context the beneficial role of short-term depression (STD) is illustrated. After summarizing heuristic indications emerging from the study performed, we conclude by briefly discussing open problems and critical issues still to be clarified.

© 2004 Elsevier Ltd. All rights reserved.

Keywords: Working memory; Learning; Synaptic plasticity; Spike timing dependent plasticity; Synaptic frequency adaptation; Spiking neurons

1. Introduction

Modelling the dynamic behavior of large populations of neurons aims primarily at capturing essential principles underlying the way in which the interaction with the environment shapes the structure and the behavior of the central nervous system.

For such a formidable task to be feasible, it has to be partitioned in sub-tasks of manageable complexity. This implies, among other things, mapping the model architecture onto a neurobiological structure of interest, identifying the input and output channels of the sub-system being modelled and making assumptions as to the coding strategy relevant in the case at hand.

The critical stage in which the model is formulated, its building blocks defined and their dynamics assigned, has many degrees of freedom, which often tend to make the model itself ‘underdetermined’ with respect to experimental data. In some cases, however, experimental indications appear to decisively suggest, if not a specific model, the ‘logic’ for its construction.

The neurophysiological experiments exposing the neural substrate of ‘working memory’, or ‘active memory’ states in the infero-temporal (IT) or pre-frontal (PF) cortices seem to constrain to a significant extent the logical requirements to be met by a successful model of those functions. A source of inspiration was found in particular in experiments in which a neural population, after responding to a familiar stimulus, was able to keep firing at higher than spontaneous rate, in a stimulus selective way, after removal of the stimulus, for macroscopically long times (see Section 2).

* Corresponding author.

E-mail addresses: paolo.delgiudice@iss.infn.it (P. Del Giudice), fusi@cns.unibe.ch (S. Fusi), mattia@iss.infn.it (M. Mattia).

It was suggested (see [1]) that a most natural way to account for such sustained reverberations is an ‘attractor’ dynamics at work in a system of interacting neurons with high feedback.¹ The seminal suggestion steered a successful research program, of which we report here the portion more related to our own work.

The stimulus selectivity of the reverberant, attractor activity, i.e. the fact that the familiarity of the stimulus triggers it, points directly to the issue of *learning*, meaning in this case the development of a visual memory, i.e. the formation of *internal representations* of visual stimuli. In such a scenario, the input to the system is the activity elicited in the neural population by the visual stimulus (however encoded by previous stages of processing), and the output is the activity pattern of that same population, after the reverberation sets in, and it is an autonomous reflection of the feedback dynamics of the system.

The long standing suggestion, that the substrate of learning is to be found in the modifications of the synaptic couplings among neurons, together with the original proposal of Hebb [2], that those modifications are *local*, and dependent on the covariance of the pre- and post-synaptic activity for each synapse, take a precise and quantitative form in the theory of attractor neural networks [3–6], and generate predictions that can be in principle (and have been in part) verified in experiments.

Recurrent neural network models incorporate *learning* through various kinds of prescriptions, most of which capture the original Hebbian suggestion, by which learning is associated with synaptic long-term potentiation (LTP) and is driven by the covariance of pre- and post-synaptic activities. The structure of incoming stimuli is transferred into the synaptic development by means of the frequencies: incoming stimuli activate (raise the frequency of) given sets of neurons, and the synaptic efficacies are pre-specified functions of the pre- and post-synaptic frequencies. The *Hebbian* facilitation mechanism for high pre- and high post-synaptic frequencies, is accompanied by some choice for the mechanism in charge of long-term synaptic depression (LTD).

These *learning rules* proved to account for a very rich phenomenology, including retrieval and maintenance of ‘learned’ internal representations of a stimulus in tasks implying working memory (see e.g. [1,5,6]).

As richer and richer collective behaviors of large networks of spiking neurons are understood and brought ‘under control’ (in terms of analytical description and/or of semi-realistic simulations), and experi-

mental findings shed some light on the biological mechanisms of LTP and LTD, it appears timely to explore ways of incorporating the *dynamics* of learning into the model.

This is not only a natural quest of a dynamic substrate for supposedly known features of learning. Instead, it opens up an entire new territory for the dynamical analysis of recurrent networks of spiking neurons.

In this paper we summarize the neurophysiological evidence related to the establishment of working memory, set up a general framework for a theoretical approach to learning, outline some general requirements to be met by a mechanism of synaptic structuring and some of the questions raised by the inclusion of synaptic dynamics, sketch possible ways to put under partial control the search for interesting dynamical regimes in the huge neural *and* synaptic parameters space, and provide examples for a specific case.

The plan of the paper is as follows: Section 2 summarizes experimental evidence for stimulus-specific, enhanced activity observed in electrophysiological recordings from various cortical areas during tasks involving a delay. We focus in particular on infero-temporal cortex (IT) (and briefly mention experiments in pre-frontal cortex, PFC), and emphasize the emergence of delay activity as an automatic consequence of the building up of long lasting archetypical neural representations of visual objects, even when it is not needed for performing successfully the task (as in the case of simple delayed matching to sample (DMS) tasks, where long-term memory is not needed).

In Section 3 The ‘attractor’ theoretical framework is summarized: the Hebbian view is implicitly adopted, of a distributed reverberation as the global outcome of local synaptic changes driven by covariant pre- and post-synaptic activities. The delay activity would be interpreted as a self-sustaining, selective attractor state ignited by the appearance of a ‘familiar’ stimulus, and kept alive by the high feedback.

Section 4 introduces a family of models of spiking neurons and of assemblies of those, interacting via excitatory and inhibitory synapses. We describe the widely used ‘integrate-and-fire’ (IF) neuron, the approach to its dynamics in terms of stochastic afferent currents, and the ‘mean field’ treatment of the collective equilibrium states for an interacting population of IF neurons. We briefly discuss how a synaptic structure affects the equilibrium properties of the system and allows for self-sustaining selective activity. We also briefly mention an approach to a dynamic mean field treatment.

In Section 5 we sketch a conceptual frame for ‘learning’ and ‘forgetting’ the suitable synaptic structure in a quite general realistic scenario, in which synapses are discrete and bounded, and the global dynamics of learning is described as a walk through the discrete synaptic state space. Considerations of the network’s

¹ This is to be understood as a logical constraint, and its power is not spoiled by the possibility that, for example, the neural area of interest might be receiving sustained input from elsewhere in the brain, since this would only shift the problem to the justification of that input signal. See the discussion in [1].

memory capacity lead to the proposal of a ‘stochastic’ (and slow) learning dynamics. From the analysis, which encompasses many possible ‘learning rules’, the Hebbian rule emerges as the optimal one in terms of signal-to-noise ratio.

In Section 6 a family of ‘microscopic’ implementations of the stochastic learning idea in terms of detailed synaptic dynamics is presented. The dynamic synapse is spike-driven, in such a way to implement a rate dependent, Hebbian long-term plasticity. Concerning the hot ongoing debate on which models of synaptic plasticity are presently best supported by experimental evidence, we remain neutral in this review on this issue, and we anticipate that the dependence of the synaptic changes on the relative timings of pre- and post-synaptic spikes and/or on the depolarization of the post-synaptic cell in the models we describe, will be instrumental for better implementing a rate-based Hebbian dynamics, having the synapse autonomously ‘measuring’ the rates on a local (in time and space) basis. We will just touch upon the biological plausibility of the specific models presented at the microscopic level.

In Section 7 we set the stage for the study of the coupled, autonomous, ongoing dynamics of neurons and synapses. We identify critical features affecting the Hebbian character of the synaptic models introduced in the previous section. A number of constraints is then listed, ensuring that the pattern of stable equilibrium states available to the system is compatible with plausible learning scenarios. To this end, a mean field approach is employed to guide the search of interesting regions in the network’s state space.

In Section 8 we provide a working example of DMS states autonomously forming in a simple stimulation protocol, as a result of the coupled dynamics of neurons and synapses, and use it as an illustration of general features of the problem, and list critical requirements for a ‘learning trajectory’ to be successfully completed.

In Section 9 we briefly mention a small subset of the many open issues, and list some simplification that should be relaxed in order to come closer to experimental evidence.

2. Delay activity in experiments

Persistent enhanced spike rates have been observed throughout the delay interval between two successive visual stimulations in several cortical areas: after the pioneering work by Fuster [7] and Niki [8], they have been encountered in infero-temporal cortex (IT) [9–15], in pre-frontal cortex (PF) [16–20] and in posterior parietal cortex [21] (for reviews see [22–24]). The phenomenon is detected in single unit extra-cellular recordings of spikes following, rather than during, the presentation of a sensory stimulus. It has been found in

primates trained to perform a delayed-response task, in which the behaving monkey must remember the identity or location of an initial eliciting stimulus in order to decide upon its behavioral response. The pattern of delay activity is considered a neural correlate of working memory, which is related to the ability of the animal to actively hold an item in memory for a short time. In fact, lesions in IT and PF are known to produce impairments in spatial and associative memory tasks [25,26], and affect recognition of an object after a delay in both humans and monkeys.

We mention for completeness that different types of stimulus-specific delay activity have been observed (e.g. in PF, see [28]) in tasks involving working memory, such that the neurons’ firing during the delay encodes a scalar quantity characterizing the stimulus. This kind of delay activity is left out in the models reviewed in the present work (see e.g. [29]).

2.1. The experimental protocols

The features of the observed delay activity depend on the area, on the experimental protocol and on the task performed by the animal. The prototypical experimental protocol used to investigate delay activity and working memory is the delayed matching to sample (DMS) task. A trial begins with a brief presentation of one visual stimulus, the sample stimulus. After a delay of a several seconds, a second image (match stimulus) is presented to the animal which has to respond differently if the second stimulus is identical to the sample or not (see e.g. [11,12]). In some cases the stimuli are compound and the monkey has to remember during the delay the individual features of the compound stimuli (e.g. color or shape) [9]. In one experiment [15], the sample stimulus is presented a random number of times, and the monkey has to release a lever as soon as the sequence of sample stimuli ends with the presentation of a different stimulus. There are also tasks in which multiple intervening stimuli are presented between the sample and the matching test stimulus [14,18].

In the experiments that involve PF recordings, the tasks are designed to engage both spatial and object memory. In the experiments of the Goldman-Rakic group [16,17,19] the animal performs an oculomotor delayed-response (ODR) task in which the location of the stimulus has to be retained during the delay. The monkey is required to fixate a point while the sample stimulus is presented. Following a GO signal, the monkey has to make a saccadic eye movement to the location of the sample stimulus. In one experiment [20], the monkey was trained to remember both the location and the identity of the stimulus. The trial starts with a sample stimulus in the center of the screen. After a first delay (a “what” delay), two objects, one of which is the sample, are presented simultaneously in two different

positions of the screen. Now the monkey has to identify the matching stimulus and to remember “where” it is during the following, last delay of the trial and to make a saccade to the remembered location of the matching object.

All these tasks require active memory, and can be performed by the monkey also when novel stimuli are presented. In [13], the authors trained the monkey to perform the paired-associate learning test that is usually adopted by clinical neuropsychologists to assess human long-term memory. In a first stage the monkeys learned a set of pairs of computer generated pictures. In each trial the sample stimulus is presented to the monkey, and, after a delay of 4 s, the paired associate of the sample and one from a different pair are presented simultaneously. The monkey is rewarded if it touches the associate. The pair-association task clearly involves long-term memory since the monkey has to retrieve the other member of the pair associated with each sample stimulus, and there is no way of performing the task by using the information available in a single trial.

2.2. Features of the observed delay activity

The recorded delay activity has different features depending on the area. The delay activity phenomenology is much richer in PF than in IT, probably because IT is the last exclusively visual area in the ventral pathway. To illustrate the typical features of the delay activity we show in Fig. 1 an example of cortical recordings in IT while the monkey was performing a DMS task. The main features of delay activity in IT are:

- The average spike frequency in the delay can be stimulus selective and is highly reproducible: each stimulus evokes a characteristic pattern of delay activity. Usually a given neuron or neural population is selective to one or a few visual stimuli (e.g. in [27] neurons with selective delay activity respond on average to 5 out of 30 stimuli when abstract geometrical patterns are used as visual stimuli). In IT there is no dependence on the position, especially in anterior IT, where neurons tend to respond to complex stimuli rather than to simple features [30].
- When stimuli are presented in a fixed temporal order, the delay activity pattern encodes in its spatial structure the temporal association between stimuli: if one neuron responds to a given stimulus (during the stimulus presentation and during the delay), it usually responds also to the neighboring stimuli in the temporal sequence [12], even if the stimuli are visually unrelated (see also Section 9).
- The typical selective rates in the delay interval are around 10–20 sp/s, against a background spontaneous activity a few spikes per second.
- The elevated rate distributions persist for very long times (compared to the inherent time scales of single cell

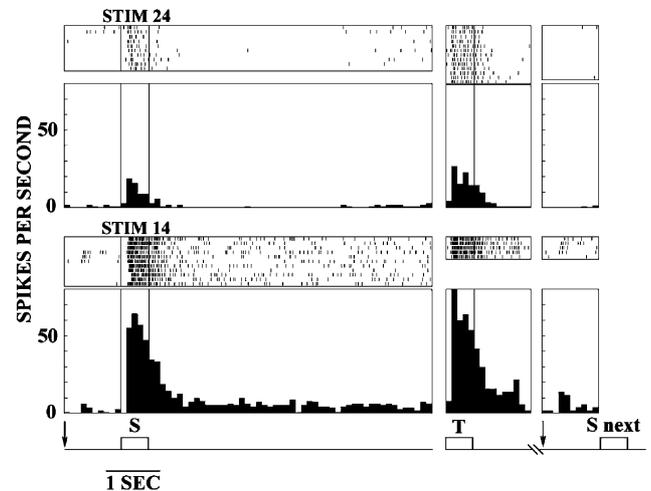


Fig. 1. Stimulus-selective delay activity in infero-temporal cortex during a DMS task. The animal is required to compare a sample stimulus to a test stimulus which is presented after a delay of a few seconds. The response of one cell to two different familiar stimuli (stimulus 24 and stimulus 14) is shown in the plot. The rasters show the spikes emitted by the cell in different trials, and the histogram represents the mean spike rate across all the repetitions of the same stimulus as a function of time. The visual sample stimulus *S* triggers a sustained delay activity in response to stimulus 14, but not to stimulus 24. The information about the last stimulus seen is propagated throughout the inter-stimulus interval, no matter whether the identity of the stimulus has to be kept in mind or not to perform the task. Indeed the elevated activity elicited by stimulus 14 is triggered also in the inter-trial interval, between the test stimulus and the sample of the next trial, where there is no need to hold in memory the identity of the last stimulus seen (adapted from [27]).

dynamics), as long as 20 s [9]. The typical delays used in the experiments are around 5 s (e.g. 6–20 s in [9], 0.5–5 s in [15], up to 16 s in [11,12], 4 s in [13], 4–8 s in [27,31]).

- The pattern of delay activity reflects the last stimulus seen. Intervening stimuli disrupt the delay activity [14].
- When unfamiliar stimuli are used (e.g. computer generated geometrical pictures), many repeated presentations of the same stimuli are needed before selective delay activity appears [32]. Unfamiliar stimuli have never been observed to evoke delay activity [12]. Hence stimulus selectivity in IT seems to be acquired through training [33] and the learning process is slow. In other areas learning might be faster, but still requires tens of presentations. For instance Erikson and Desimone report in [34] that reaching a performance of 85% in a pair-associate task requires on average 30 presentations per pair. Whether this level of performance is related to the formation of delay activity somewhere in the cortex is still unclear.
- There is no correlation between the erroneous response and the absence of delay activity in DMS tasks. In [15] the authors compared the firing rate during the final delay period in correct trials, to the one during the delay period just before the erroneous response and no

clear difference was detected. Another indication that delay activity is not necessary to perform the DMS task comes from the previous point in this list: no delay activity was observed for novel stimuli but the monkey was able to perform well anyway. These considerations do not rule out the hypothesis that the delay activity observed in IT is actually used to compare the sample and the test stimulus. The only conclusion that can be drawn is that some other mechanism, probably located in other areas, is certainly involved to perform the DMS tasks.

- Delay activity is mechanistic: in DMS tasks it is automatically evoked by any familiar visual stimulus, including the test stimulus that has not to be retained in memory to perform the task. Hence delay activity is not linked by the task and the information about the identity of the test stimulus can be propagated up to the presentation of the sample stimulus of the next trial, surviving the reward phase and long inter-trial intervals in which the monkey is not fixating the screen [27]. This phenomenology suggests a possible role of delay activity in IT: it is the best candidate to preserve actively the memory of events that are separated in time and hence it might provide a simple mean by which temporal sequences and context dependent memories can be encoded in IT [12,27].

- In one of the experiments [15], the delay activity evoked by partial stimuli was studied. The firing elicited by a limited portion of the sample stimulus was comparable with that elicited by the original sample stimulus. It is not described what happens to the visual response during the presentation of the partial stimulus. In another experiment [31], when familiar visual stimuli are progressively degraded by superposing Gaussian noise to the visual stimulus, the visual response changes gradually while the evoked delay activity remains unchanged up to some level of degradation and then suddenly disappears for practically unrecognizable visual stimuli.

- In a few cases elevated delay activity was observed also when the sample stimulus did not elicit any visual response [13].

- The percentage of recorded cells that exhibited selective delay activity depends on the number of stimuli presented and is usually low. Furthermore those cells tended to be localized in remarkably restricted areas [32].

The phenomenology in PF cortex is richer and more complex. Relevant features of the recorded cells exhibiting delay activity are:

- The delay activity is selective either for position or for objects. In [20] it is shown that there are cells which are selective for both. When spatial information is encoded, a continuous spectrum of patterns of localized persistent activity is observed. Each pattern represents a specific position which is an inherently analog variable (see [24] for a review).

- The sample-selective delay activity is maintained throughout the trial even when other test stimuli intervened during the delay [18].

- The activity observed in the delay is present also in the inter-trial interval and it is not related to eye movements or to any other observable behavior that followed the test stimulus [19].

- The delay activity is either truncated or totally absent on erroneous trials in experiments in which the monkey has to perform an oculomotor delayed-response task [16].

- The percentage of cells with selective delay activity is higher than in IT, but the duration of the persistent activity might be shorter. In [19] the monkey has just to maintain fixation of the visual stimulus and the delay activity lasted up to 2.6 s. It is not clear in the other cases because the delays are usually shorter than 2 s. In one experiment [16] the monkey performed an oculomotor delayed-response task with delays up to 6 s and the delay activity was observed to persist throughout the whole delay interval.

3. The attractor picture

The experimental findings of the DMS experiments have been interpreted as an expression of an attractor dynamics in the cortical module: a comprehensive picture has been suggested which connects the pattern of delay activity to the recall of memory into an active state [1]. It is mainly inspired by the results of the DMS experiments in IT, but the same principle can be probably exploited for explaining the entire delay activity phenomenology described in the previous section.

The collective character is expressed in the mechanism by which a stimulus that had been previously learnt, has left a synaptic *engram* of potentiated excitatory synapses connecting the cells driven by the stimulus. When this subset of cells is re-activated by a sensory event, they cooperate to maintain elevated firing rates, via the same set of potentiated synapses, after the stimulus is removed. In this way, the cells in each group can provide each other, with an afferent signal that clearly differentiates the members of the group from other cells in the same cortical module.

The collective nature of the pattern of delay activity is related to its attractor property. Since many cells cooperate, the pattern of delay activity is robust to stimulus “errors”, i.e. even if a few cells belonging to the self-maintaining group are absent from the group initially driven by a particular test stimulus, or if some of them are driven at the “wrong” firing rates, once the stimulus is removed, the synaptic structure will reconstruct the ‘nearest’ distribution among elevated activities of those it had learned to sustain. The reconstruction (the attraction) will succeed provided the deviation from

the learned template is not too large. All stimuli leading to the same pattern of delay activity are in the same *basin of attraction*. Each of the stimuli repeatedly presented during training creates an attractor with its own basin of attraction (see e.g. [1,4,5]). In addition, the same module can have all neurons at spontaneous activity levels, if the stimulus has driven too few of the neurons belonging to any of the stimuli previously learned.

Whether this attractor dynamics is localized in a specific area or it requires complex excitatory circuits that involve different cortical or sub-cortical (e.g. thalamic) areas is still debated. The delay activity in a given area might be a mere reflection of the persistent activity localized in a different area and might not necessarily take part in the attractor dynamics. However there is accumulating evidence that the delay activity in IT is affected by top-down inputs coming e.g. from PF cortex [35]. Less direct evidence is available for interactions through the sub-cortical pathway [23]. It seems a reasonable working hypotheses to assume that the principles underlying attractor dynamics do not change much if the delay activity is the result of the interaction of several cortical modules.

3.1. Experimental evidence for the attractor picture

In the framework outlined above the observed delay activity is not a single neuron property but a result of the interaction of a large number of neurons. For IT two strong arguments support this view:

- When the delay activity is stimulus selective there is a strong positive correlation between the visual response and the delay activity in the following inter-stimuli interval, when the average activity for each stimulus is considered. However, the correlation between the activity during the presentation of the sample stimulus eliciting the maximal response in the delay and the inter-stimuli interval delay activity on a trial-by-trial basis is much weaker. In fact, the visual response and the activity in the last second of the inter-stimuli interval are generally uncorrelated [27,31]. Although the delay activity immediately following a specific stimulus depends on the magnitude of the visual response, which can vary from trial to trial for different reasons, the final level of delay activity (a few seconds later) is constant. This supports the suggestion that the delay activity is a result of the neural network properties, rather than a change in the state of the single neuron alone, triggered by the visual response.
- When the visual stimulus is degraded [31] or partially obscured [15], the internal representation during the stimulation can vary dramatically but the delay activity response remains unchanged provided that the vi-

sual stimulus is still recognizable. This strongly supports the idea that the delay activity pattern is the result of the cooperation of a large number of interacting different cells.

4. Dynamics of networks of model neurons with fixed synapses

In this section we (i) introduce the family of ‘integrate-and-fire’ (IF) neurons, workhorses of a large part of modelling studies so far, (ii) briefly describe the dynamics of IF neurons with randomly fluctuating input, (iii) formulate the ‘mean field’ approach to the equilibrium states of a population of interacting IF neurons, (iv) define a prototype network architecture involving excitatory and inhibitory synaptic couplings, and selectively long-term potentiated and depressed synapses.

Most of the discussion will be based on the approach initiated in [5,36,37].

We briefly remind the essential features of the model neuron (for extensive reviews see [38,39]). IF neurons are one-compartment, point-like models, and the state of the neuron is completely described at each time by the instantaneous value of its membrane depolarization V . The time evolution of V is given by

$$C\dot{V} = L(V, t) + I(t) \quad (1)$$

where $L(V, t)$ is a leakage term embodying the action of the ion channels restoring the rest membrane potential in the absence of afferent current and $I(t)$ is the total afferent current to the neuron. Frequently used forms for $L(V, t)$ include the $L(V, t) = -V/R$ (*leaky IF neuron*, R is the membrane resistance and $\tau = RC$ is the membrane time constant) and $L(V, t) = -\beta$ (*linear IF neuron*, β constant). The two forms produce similar collective behavior if a rigid lower bound to the depolarization of the linear IF neuron is imposed [40]. The IF neuron dynamics is complemented in both cases by the condition that when V reaches the emission threshold θ a spike is generated, V is reset to a prescribed value H and further spikes emission is prevented for the duration of the absolute refractory period τ_r , no matter how strong $I(t)$. The afferent current $I(t)$ has an external and a recurrent component. All neurons in the network get external excitatory input, to be interpreted as a non-specific background external activity or serving as a coding of incoming external stimuli. Spikes continuously exchanged among the neurons inside the network constitute the recurrent part of the current; in the simplest case the recurrent afferent current is described as a series of impulses, each one depolarizing (EPSP—excitatory post-synaptic potential) or hyperpolarizing (IPSP—inhibitory post-synaptic potential) the membrane of the receiving neuron depending on the excit-

atory or inhibitory nature of the emitting neuron. We consider only fast synaptic conductances (AMPA mediated for excitatory neurons and GABA_A mediated for inhibitory neurons) and we assume that the details of the conductance dynamics can be neglected. This is usually a good approximation for fast synaptic conductances. Slow synaptic conductances (NMDA) might play an important role in making the selective delay activity more stable and they have been hypothesized to be a fundamental element when the attractor corresponding to a given stimulus survives the presentation of distractors in PF [18,41]. See [24] for a synthetic explanation of the basic mechanism and for a review of papers that discuss the role of slow synaptic currents.

The *efficacy* of the synapse connecting two given neurons equals the amount of EPSP or IPSP provoked by the spikes transmitted. Consistently with prevailing biological evidence, synaptic efficacies are assumed to be a small fraction of the allowed range of values for V . In the simplest case the total afferent current due to the spikes emitted by pre-synaptic neurons is:

$$I(t) = \sum_j J_j \sum_k \delta(t - t_j^{(k)} - d_j^{(k)}) \quad (2)$$

where J is the synaptic efficacy, j labels the pre-synaptic neurons, $t_j^{(k)}$ is the time the k th spike is emitted by the j th pre-synaptic neuron and $d_j^{(k)}$ the corresponding transmission delay. All the types of synaptic couplings EE, II, EI, IE are present, EE being the only plastic ones (E: excitatory, I: inhibitory).

Due to the various sources of irregularities in the network in typical conditions, the total afferent current can be described as a stochastic, Gaussian driving signal, such that $V(t)$ is itself a Gaussian stochastic process. Then, a *transfer function* can be defined and computed for the neuron, giving the mean output neuron frequency ν_{out} as a function of the mean μ and variance σ^2 of the afferent current $I(t)$. This function can be computed analytically for both the linear [40] and the leaky IF neurons [42]. In the case of the linear IF neuron Φ has a particularly simple form:

$$\nu = \left[\tau_f + \frac{\sigma^2}{2\mu^2} \left(e^{-\frac{2\mu\theta}{\sigma^2}} - e^{-\frac{2\mu H}{\sigma^2}} \right) + \frac{\theta - H}{\mu} \right]^{-1} \quad (3)$$

Interestingly this function contains all the single neuron properties that are relevant when the neuron is embedded in a network of interacting cells [36,43]. Such a function has been measured for cortical cells in vitro and turned out to be well described by the theoretical transfer function of both IF models [44].

The pattern of connectivity is chosen to be sparse (each neuron receives spikes from a small fraction of the other neurons in the network) and random, but fixed (no geometry is imposed on the connectivity pattern, which

is chosen at the beginning by performing for each neuron a random selection of its pre-synaptic contacts).

Sparse connectivity (small probability of shared input), together with small synaptic efficacies, allows to assume that the spikes impinging on each neuron are many and statistically *independent* events, and a *mean field approach* can be formulated.

In the approximated, ‘mean field’ (MF) approach, the network is partitioned in ‘populations’ of neurons, all neurons in a population sharing the same statistical properties of the afferent currents, and firing spikes *independently* at the same rate. Fig. 2 illustrates the architecture of the model.

In this theoretical framework, the single neuron transfer function (providing the output rate as a function of the average input current) is turned into a ‘population transfer function’ Φ : μ and σ^2 [5] are then expressed in terms of the rates of the corresponding populations (including external neurons), and the set of stationary, self-reproducing rates for the different populations are found by solving a set of coupled

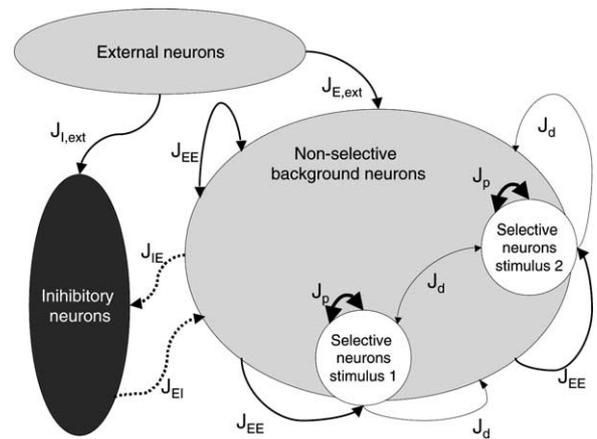


Fig. 2. Schematic illustration of the network’s architecture. Blobs represent different populations in the network: Inhibitory (I) neurons (dark gray blob), excitatory (E) neurons which are never stimulated (‘background’, light gray blobs), neurons that, because of the formation of a stimulus-specific synaptic structure, are selective to different stimuli (two stimuli in the figure; white blobs). The sets of selective neurons are assumed to be disjoint in the figure, being the associated stimuli uncorrelated. In the general case selective blobs overlap, owing to neurons shared by internal representations of different stimuli. Thick arrows indicate potentiated synapses (J_p), connecting neurons coding for the same stimulus. Medium weight arrows indicate synapses which do not undergo modifications with respect to the initial state (J_{EE}), connecting background neurons among themselves, and background neurons to selective ones. Thin arrows stand for depressed synapses (J_d), assuming a *homosynaptic* kind of LTD (i.e. synapses get depressed for highly active pre-synaptic neurons and low activity post-synaptic ones); they connect selective neurons to the background, and neurons selective to different stimuli to each other. Synapses from (J_{EI}) and to (J_{IE}) inhibitory neurons are depicted as dashed lines. External neurons interact with the network via the synaptic couplings $J_{I,ext}$ and $J_{E,ext}$, and are assumed to be excitatory.

self-consistency equations, which for a single population reduces to

$$v = \Phi(\mu(v), \sigma^2(v)) \equiv \Phi(v). \quad (4)$$

In practice, to solve the above fixed point equation(s) one generates a fictitious dynamics which generates an iterative process that converges to the solution. A typical choice for the fictitious dynamics is one following the gradient of Φ .

Fig. 3 illustrates how the solutions of Eq. (4) depend on the average synaptic coupling in a given sub-population: starting from a low-coupling situation in which the only available fixed point is one of low rate, spontaneous activity, increasing (excitatory) couplings bend the transfer function, up to the point in which it gains two more solutions of the self-consistency equations (4). The upper fixed point corresponds to the selective state with enhanced rate.

Given a fixed point v^* solving Eq. (4), it is relevant to assess its stability, i.e. its property of being restored following a small perturbation.

If one could solve the dynamic equation for the system, the explicit time-dependent solution would be used of course to assess the stability. This is often not the case, and one has to resort to approximated strategies. One popular choice, to be checked a posteriori through simulations, is to use the fictitious dynamics introduced above to assess the system's stability.

Following this approach it is easy to conclude that the white circles in Fig. 3 correspond to stable fixed points, while the middle fixed point is unstable. $\Phi' < 1$ is the stability condition for the excitatory population. It has also been recognized in several works [45–50] that a proper treatment of the network behavior near the fixed point requires the formulation of the actual (mean field) dynamics. In particular it was shown in [50] how $\Phi'(v^*)$

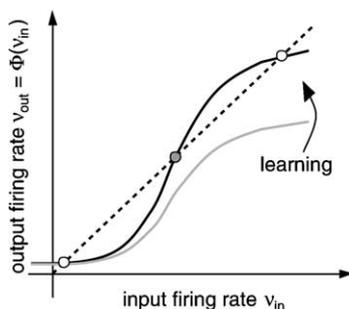


Fig. 3. Population transfer functions $\Phi(v_{in}) (= v_{out})$ and fixed points (self-reproducing rate states) of a stimulus-selective neuron subset before and after learning. Solid curves are $\Phi(v_{in})$ before (light gray) and after (dark gray) the recurrent coupling potentiation: When J_p increases the same v_{in} drives more excitatory current, amplifying the output spike emission rate v_{out} . Fixed points are such that $v_{out} = v_{in}$ (the self-consistency equation), the intersection of v_{in} (dashed line) and the transfer functions: White circles are for stable states, dark circle for the unstable one.

enters in a simple way the description of the dynamics near a fixed point for a single neural population (which is a good approximation for low enough coding level f).

For a given network's architecture, and number and coding level of afferent stimuli, the set of rates solving the MF equations are functions of the average synaptic efficacies, depending on which we might have only low rate, 'spontaneous' states for all the populations (the lowest white circle in Fig. 3), or those might coexist with higher rate, stimulus selective states (the highest white circle).²

We can depict the underlying dynamical scenario as an essentially quiescent network in a 'spontaneous' state (all neurons firing at very low rates) in the absence of stimulation, then the extra current pumped by an incoming stimulus temporarily makes a sub-population of neurons increase their rates by a large amount; after releasing the stimulus, depending on the structure of the synaptic couplings, the network can relax back to the spontaneous state, if this is the only one available, covering with its basin of attraction the whole network state space; or, if a stimulus selective, collective state is also available as an attractor of the network's dynamics, the stimulation just released might have pushed the network's state over the barrier separating the spontaneous and the selective attractor, and the network ends up in the selective state. In Fig. 4 is shown how a simulation of a network of IF neurons with structured synaptic couplings reproduces such a scenario.

Another possible implication of the potentiation of excitatory couplings is suggested in [50]: the more potentiated the excitatory synapses, the faster the network reaction to a stimulus. Fig. 5 shows the theoretical prediction [50] for the transient response of an excitatory population which, starting from a spontaneous stable state of low emission rate, undergoes a sudden jump in its external input, providing an high emission frequency response. For the same initial and final asymptotic average emission rate, higher average excitatory couplings entail quicker response, and faster damping of the oscillations. If confirmed in more realistic scenarios, this effect could be related to experimental indications of a reduction in the visual response latency following training [34].

5. Realistic learning prescriptions

The suitable synaptic structure for sustaining a large number of patterns of delay activity is learnt automatically and is likely to be an expression of the synaptic dynamics. Before going into the details of a realistic

² For the many population case, there have been attempts to derive an effective, single-population transfer function [51].

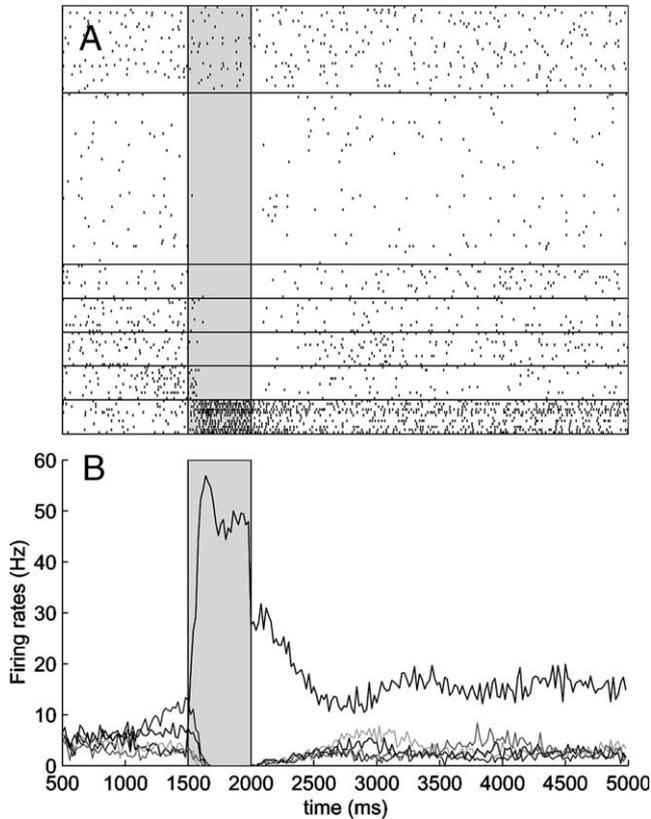


Fig. 4. Stimulation of persistent delay activity selective for a familiar stimulus in a simulation of interacting spiking neurons with synaptic coupling structured as in Fig. 2. The raster plot (A) shows the spikes emitted by a subset of cell in the simulated network grouped in sub-populations with homogeneous functional and structural properties: The bottom five strips contain the spike traces of neurons selective to five different uncorrelated stimuli; In the upper strip is shown the spike activity of a subset of inhibitory neurons and in the large middle strip several background excitatory neurons. (B) The emission rates of the selective sub-populations are plotted: The activity of a population is given by the fraction of neurons emitting a spike per unit time. The population activity is such that before the stimulation all the excitatory neurons emit spikes at low rate (global spontaneous activity) almost independently of the belonging functional group.

synaptic dynamics we consider a wide class of biologically plausible prescriptions for updating the synaptic efficacies. The general conclusions we will be able to draw from a very limited number of reasonable and simple assumptions will provide the guidelines for designing a detailed model of synaptic dynamics (see [52] for a more detailed review).

During the presentation of each stimulus, a pattern of activities is imposed onto the network and each synapse is updated to encode the information carried by the pre- and the post-synaptic neuron. It is reasonable to assume that each synapse sees only the activities of the two neurons it connects (locality in space) and that the inter-stimuli intervals can be so long that the activity read by the synapse does not depend on the past history and hence that the updating can be performed only on the

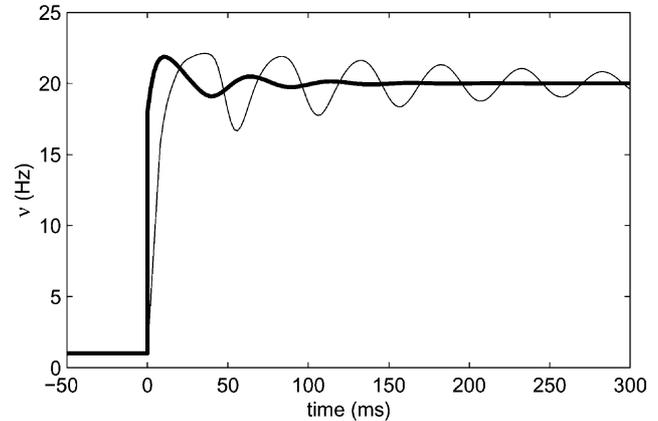


Fig. 5. Transient responses to a stepwise stimulation of an excitatory population, for two different coupling strengths. The theoretically predicted spike emission rate dynamics of the population with weak (thin line) and strong (thick line) synaptic coupling is shown (adapted from [50]).

basis of the stimulus currently presented (locality in time). As for the structure of the synapse, it is reasonable to assume in general that all the internal variables describing the synaptic state are bounded and that long-term modifications of the synaptic internal variables cannot be arbitrarily small, regardless of the details of the synaptic model.

5.1. The palimpsest property: a tight constraint on storage capacity

Under the above general assumptions, any network of neurons exhibits the ‘palimpsest’ property [53–56]: old stimuli are automatically forgotten to make room for the most recent ones. The memory span is limited to a sliding window containing a certain number of stimuli that induced synaptic modifications. Within this window, recent stimuli are best remembered while stimuli outside the memory window are completely forgotten. The width of the sliding window depends on how many synapses are changed following each presentation: if this number is small the patterns (stimuli) to be learnt have to be presented many times (slow learning) but eventually the memory resources are equally distributed among the different stimuli.

If the number of synapses that are modified upon each presentation is large the network learns quickly but the memory span can be so limited that it might compromise the functioning of the network as an associative memory [52,56]. In fact it can be proven that the maximum number p of random uncorrelated patterns that can be stored in the synaptic matrix scales at most as:

$$p_{\max} = k(r, s) \log N$$

where N is the number of neurons in the network and k is a constant (i.e. not dependent on N) that depends in a

complicated way on r , i.e. the ratio between the minimal long-term modification that can be induced and the full range in which the internal synaptic variable varies, and on s , i.e. the mean fraction of synapses that are changed upon each stimulus presentation. r depends on the inherent synaptic dynamics, while s depends on the statistical structure of the patterns to be encoded (for more details see [52]). $k(r, s)$ is a monotonic function of r and s and tends to infinity when r or/and s go to 0 in the limit $N \rightarrow \infty$.

This constraint is extremely tight and very general and makes the network a very inefficient associative memory. If one of the parameters on which r or s depend becomes N -dependent then it is possible to extend dramatically the memory span. For instance if the number of possible synaptic states increases with N (as in the case of the Hopfield model [4]), k provides an extra N -dependent factors which in some cases destroy the palimpsest behavior. In those cases the storage capacity is mainly limited by the interference between the memory traces of different stimuli, and not by memory fading. For a wide class of models, as soon as the maximum storage capacity is surpassed, the network suddenly becomes unable to retrieve any of the memorized patterns. This is also known as the blackout catastrophe (see e.g. [57]) and in order to prevent it, one has to stop learning at the right time, or the network must be able to forget [53]. Hence the palimpsest property, if forgetting is not too fast, can be a desirable feature of the learning process.

5.2. Back to the optimal storage capacity: the stochastic selection mechanism

Decreasing the fraction of synapses that are changed upon each presentation can increase dramatically the memory span. Depending on the statistical properties of the patterns, the network can acquire enough information about them by updating only a fraction of the synapses eligible for a change according to the learning rule. What the network needs is a mechanism that selects which synaptic efficacies have to be changed following each stimulation. In the absence of an external supervisor that could perform this task, a local, unbiased mechanism could be stochastic learning: at parity of conditions (activities of the pre- and post-synaptic neurons) the long-term changes are induced with some probability. This means that only a fraction a synapses, randomly chosen among all eligible connections, are changed upon the presentation of a stimulus. In this case the number of internal states of the synapse that are stable on long time scales can be reduced to the extreme, and even networks of neurons connected by bistable synapses can perform well as associative memories. To illustrate how stochastic learning works we will study the simple case of bistable synapses in the following section.

5.3. The three basic ingredients of learning rules for auto-associative memories

In the case of a bistable (or a multi-stable) synapse the learning process can be seen as a random walk among the stable states. This kind of stochastic learning has been studied analytically and extensively in [55,56,58–60]. We assume here that the synapse has only two stable states, corresponding to two efficacies: $J^1 = 0, J^2 = J$. Each stimulus selects randomly a mean fraction f of neurons and drives them to a state of elevated activity, the same for all the active neurons. In order to estimate the retrievable memory trace of the oldest stimulus presented (ξ^1) we introduce the classical signal-to-noise ratio (see e.g. [57]): the signal S expresses the distance between the distribution of the total synaptic input across all neurons that should be active, and the corresponding distribution across the neurons that should stay quiescent when the pattern of activity ξ^1 is imposed to the network. Quantitatively S is defined as the difference between the averages of these two distributions. The noise R represents the mean width of the two distributions. High S/R would allow the network to retrieve from memory the pattern of activity ξ^1 that is embedded in the synaptic matrix and to make it a stable point of the collective dynamics (see e.g. [57]).

The synaptic updating rule is as follows: assuming for the sake of clarity that the neurons can only be in an ‘active’ (high spike rate, 40–60 Hz) or ‘inactive’ state (spontaneous firing rate, 1–5 Hz), when the two neurons connected by the synapse are both active, the synapse makes a transition to the potentiated state with probability q_{AA}^P and a transition to the depressed state with probability q_{AA}^D . The other possible transitions occur with probabilities denoted with a similar notation (e.g. q_{IA}^D is the probability that depression occurs when the pre-synaptic neuron is active and the post-synaptic neuron is inactive).

For such a system, after the presentation of p patterns, the signal corresponding to the oldest (and hence with weakest memory trace) pattern is [52,56]:

$$S = Jf\lambda^{p-1}[(1 - c_0)(q_{AA}^P - q_{IA}^P) + c_0(q_{IA}^D - q_{AA}^D)] \quad (5)$$

where c_0 is the initial fraction of potentiated synapses and λ is:

$$\lambda = 1 - (q_{AA}^P + q_{AA}^D)f^2 - (q_{IA}^P + q_{IA}^D + q_{AI}^P + q_{AI}^D)f(1 - f) - (q_{II}^P + q_{II}^D)(1 - f)^2 \quad (6)$$

which is essentially 1 minus the sum of all the transition probabilities, each multiplied by the corresponding probability of occurrence of a specific pair of activities (e.g. the probability of having that both pre- and post-synaptic neurons are active is f^2). The interference noise term depends mostly on c_∞ , the asymptotic fraction of

potentiated synapses that one would get after an infinite number of presentations of different stimuli [56]:

$$R \sim J \sqrt{c_\infty \frac{f}{N}}$$

where N is the number of neurons in the network. From the analysis of the expression of the signal-to-noise ratio one can identify at least three ingredients that characterize in general this class of updating rules:

- (1) The palimpsest property expressed by the power-law decay of the memory trace (the λ^{p-1} factor in the signal): new stimuli use resources that had been previously allocated for other stimuli, and, in doing so, erase the memory trace of the oldest stimuli. The forgetting rate depends on the statistics of the stimuli and on the transition probabilities of the synapse. In particular the network forgets fast when learning is quick (large transition probabilities, and hence, small λ s) and when the patterns are highly overlapping (high f). The memory span increases when learning is slow and the patterns are weakly overlapping.
- (2) The memory trace left by each stimulus presentation (the signal without the memory decay factor): it represents the way the synapse encodes the activity of the pre- and the post-synaptic neuron. It depends also on the initial distribution c_0 of the synaptic states. If most of the synapses are close to saturation, then the main contribution to the signal comes from those synapses which are depressed when the pre-synaptic neuron is active and the post-synaptic neuron is inactive (q_{IA}^D). Otherwise, the Hebbian term (q_{AA}^P) dominates.
- (3) The interference between memory traces corresponding to different stimuli (the noise R): the simultaneous presence in the synaptic matrix of memory traces corresponding to different patterns within the memory span generates noise that might prevent the system from recalling correctly the stored patterns. For uncorrelated patterns, when the transition probabilities and the fraction f of active neurons are small, the noise goes to zero as $1/\sqrt{N}$ and is proportional to the asymptotic fraction of potentiated synapses c_∞ , which in turn depends on the balance between the mean number of upwards and downwards transitions.

From the final expression of the signal-to-noise ratio it is clear that the most efficient way of storing patterns of activities is obtained when the Hebbian term q_{AA}^P dominates over q_{IA}^D , the transition probabilities are low ($\lambda \sim 1$) and the coding level f of the stimuli tends to zero with N . The transition probabilities should scale in such a way that all the terms in λ tend to zero with f at the same rate as $q_{AA}^P f^2$ (e.g. $q_{IA}^D \sim q_{AA}^P f$) or faster. This

would correspond to a scenario in which learning is slow—i.e. the stimuli have to be repeatedly presented to the network in order to be learnt—and the updating rule ensures the balance between the mean number of potentiations and the mean number of depressions. In such a case the network performs extremely well as an associative memory (e.g. it recovers the optimal storage capacity in terms of information that can be stored in the synaptic matrix), even if the synaptic efficacy is bistable. The number of different patterns that can be stored and retrieved from memory without errors can be as large as $N^2/(\log N)^2$, if the mean fraction f of active neurons scales as $\log N/N$ [56]. Slow learning allows also automatic prototype extraction from a class of stimuli that evoke similar, correlated patterns of activity [60,61].

6. Synaptic dynamics

Stochastic learning rules solve the problem of fast forgetting and allows for an efficient storing of the information about the statistical structure of the patterns. The stimuli can be retrieved without errors for a wide class of neural dynamics and this guarantees the stability of the patterns of activity that represent the delay activity. However, to proceed further, it is important to have a detailed model that provides a link between the transition probabilities and the synaptic dynamics driven by the neural activities. This implies the identification of the source of noise which drives the stochastic selection mechanism. The solution we propose is suggested by the analysis of cortical recordings: the spike trains recorded in vivo are highly irregular (see e.g. [62]) and the inter-spike interval variability can be exploited by the synapse to achieve stochastic transitions in a fully deterministic synaptic model.

6.1. The threshold mechanism

We now introduce a class of spike-driven synaptic dynamics that allow for stochastic learning. The most important ingredient is the existence of a threshold mechanism for consolidating a memory change only in a fraction of the cases in which the synapse would be eligible for long-term modification on the basis of the pre- and post-synaptic activity. For the sake of simplicity, we assume that the synaptic dynamics is described in terms of a single internal variable $X(t)$. If there is more than one variable describing the synaptic dynamics, then what follows applies to at least one of the internal variables, i.e. there must be at least one threshold mechanism which decides when the changes of any other variable become permanent. X is restricted to the interval $[0, 1]$ and, inside this interval, obeys:

$$\frac{dX(t)}{dt} = -\alpha\Theta(-X(t) + \theta_X) + \beta\Theta(X(t) - \theta_X) + T(t), \quad (7)$$

where Θ is the Heavyside function (it is 1 when the argument is positive, 0 otherwise). The first two terms make the upper and the lower bound the only two stable states of X when $T(t)$, the stimulus driven learning term, is 0. The internal state decays at a rate α to the lower bound $X = 0$ when $X < \theta_X$. Otherwise X is attracted towards the upper bound $X = 1$ with a constant drift β .

The importance of an internal threshold θ_X for the synaptic dynamics is at least twofold: it stabilizes memory by ignoring all the fluctuations that do not drive the internal state across the threshold and provides a simple mechanism that allows the neural activity to select which synapses are to be modified as we will see in the next section. This mechanism is so simple and robust that it can be readily implemented in analog VLSI [63,64]. The same, simple VLSI recurrent networks proved able to produce, even for deterministic external currents, highly irregular patterns of activity, suited to sub-serve a stochastic synaptic dynamics [64–66].

6.2. The learning term

The learning term determines in which direction the synapse is modified and must depend on dynamic variables related to the mean rates (the quantities to be encoded and stored in the synaptic matrix) of the two neurons connected by the synapse. We assume that the mean frequency produced by the stimulus contains all the needed information. Here we discuss two possible alternatives that have been investigated in the past years.

6.2.1. Dependence on the post-synaptic depolarization

The most classical recipe to induce LTP is the following: the post-synaptic neuron is strongly depolarized and a current is injected in the pre-synaptic neuron, causing it to emit a high frequency burst (see e.g. [67]). The protocols for inducing LTD are diverse and controversial but, again, they depend on how much the post-synaptic neuron is depolarized (a good review can be found in [68]). Here we introduce a proper dependence on the post-synaptic depolarization to achieve the scheme of modifications summarized in Section 5.3. Each pre-synaptic spike triggers a temporary modification in the synaptic variable if the synaptic threshold is not crossed. The direction of the modification depends on the instantaneous depolarization of the post-synaptic neuron $V_{\text{post}}(t)$. The jump is upwards ($X \rightarrow X + a$) if the depolarization is high ($V_{\text{post}} > V_H$), downwards if the depolarization is low ($V_{\text{post}} < V_L$). As we will see in Section 6.2.2 the distribution of the sub-threshold depolarization determines the mean firing frequency of the post-synaptic neuron and the probability of finding

the post-synaptic neuron in a high depolarization state is in general a monotonic function of the spike frequency. Two examples of synaptic dynamics under stimulation are illustrated in Fig. 6. Such a model synapse has been introduced in [63] and then studied in [69].

6.2.2. Encoding mean rates by reading depolarization

The depolarization of the post-synaptic neuron is indirectly related to the spike activity. Actually a single reading of the instantaneous depolarization does not contain much information about the post-synaptic mean rate. However the required information is distributed across several neurons that are driven to the same activity by the stimulus. This means that even a single instantaneous reading of the depolarization of a population of cells contains all the information about the mean spike rate and many other statistical properties of the activity of these neurons.

The probabilities of occurrence Q_a and Q_b of the temporary up and down regulations respectively control the direction in which the synapse is modified by the activity of the pre- and post-synaptic neurons. These probabilities depend on the statistics of the depolarization of the post-synaptic neuron under stimulation and can be calculated analytically when the model of the neuron is simple. Here we focus on the simple model of an integrate-and-fire neuron [40] introduced in Section 4. If such a neuron is injected a Gaussian current characterized by its mean μ and its variance σ^2 , the stationary distribution of the depolarization $p(v)$ has a simple expression and is given by [40,52,63]:

$$p(v) = \frac{v}{\mu} \left[\Theta(v - H) \left(1 - e^{-\frac{2\mu}{\sigma^2}(\theta - v)} \right) + \Theta(H - v) \left(e^{-\frac{2\mu}{\sigma^2}H} - e^{-\frac{2\mu}{\sigma^2}\theta} \right) e^{\frac{2\mu}{\sigma^2}v} \right]$$

where v is the mean firing frequency given by Eq. (3). Q_a and Q_b are given by the integral of $p(v)$ in the interval $[V_H, \theta]$ and $[0, V_L]$ respectively. It is straightforward to compute these integrals analytically.

μ and σ characterize the synaptic input and depend on the network interactions. We assume that v_{post} is changed by increasing or decreasing the average spike frequency of a sub-population of pre-synaptic neurons [63]. If the recurrent feedback of the post-synaptic neurons does not affect much the network activity, then the parameters of the input current move along a linear trajectory in the (μ, σ^2) space. We chose μ as an independent parameter, and $\sigma^2 = J\mu + K$. In a network of excitatory and inhibitory neurons, for which in a spontaneous activity state the recurrent input is chosen to be as large as the external input, we have that $J = J_{\text{EE}}$ (the average coupling between excitatory neurons) and $K = v_0^1 N_I J_{\text{EI}} (J_{\text{EI}} + J_{\text{EE}})$, where v_0 is the spontaneous activity of the N_I inhibitory neurons that are projecting to the post-synaptic cell (mean coupling J_{EI}). Q_a is

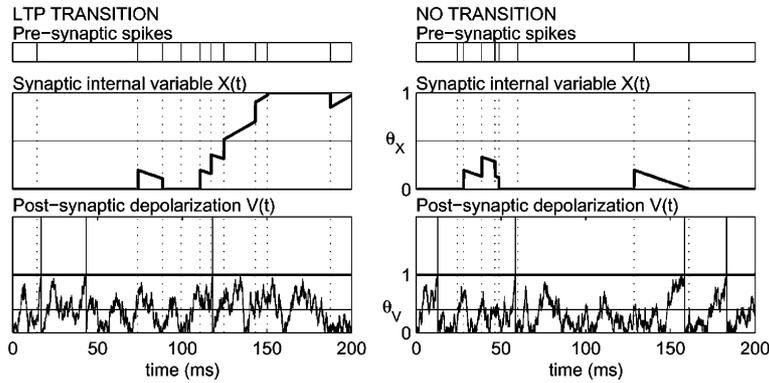


Fig. 6. Stochastic LTP: pre- and post-synaptic neurons fire at the same mean rate and the synapse starts from the same initial value ($X(0) = 0$) in both cases illustrated in the left and the right panel. In each panel are plotted as a function of time (from top to bottom): the pre-synaptic spikes, the simulated synaptic internal variable $X(t)$ and the depolarization $V(t)$ of post-synaptic neuron (note that $V_L = V_H$). Left: LTP is caused by a burst of pre-synaptic spikes that drives $X(t)$ above the synaptic threshold; Right: at end of stimulation, X returns to the initial value. At parity of activity, the final state is different in the two cases (the figure is reproduced from [63]).

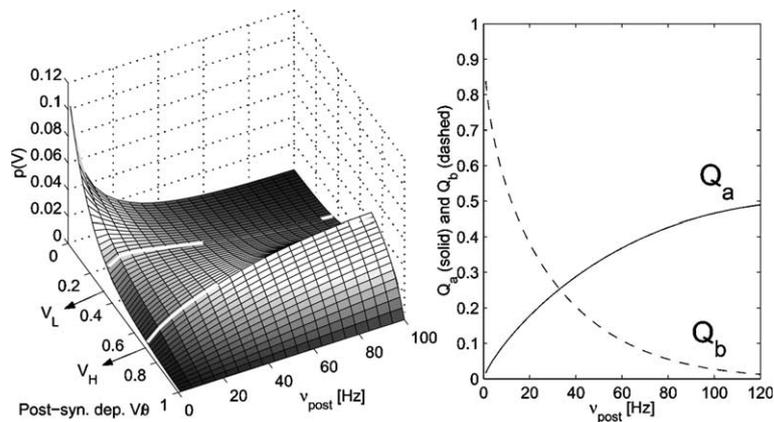


Fig. 7. Mean spike rate vs the distribution of the depolarization V . Left: Distribution $p(v)$ as a function of the post-synaptic frequency v_{post} . For each v_{post} the parameters μ, σ characterizing the input current are calculated as explained in the text. μ, σ determine the sub-threshold distribution of the depolarization that is plotted here. The white lines over the surface are drawn in correspondence of $V = V_L$ and $V = V_H$, the thresholds that determine the direction of the temporary synaptic modifications. Note that here the reset potential H coincides with V_H . Right: Probability of occurrence of upwards (Q_a) and downwards (Q_b) jumps upon the arrival of a pre-synaptic spike for different post-synaptic activities. Q_a is the integral of $p(v)$ between the white line corresponding to V_H and the threshold. Analogously, Q_b is the integral of $p(v)$ between the resting potential ($V = 0$) and V_L . As the post-synaptic activity increases, the peak of the distribution $p(v)$ moves from the resting potential to the reset potential H . As a consequence, Q_a increases and Q_b decreases (the figure is reproduced from [70]).

plotted in Fig. 7. As the external stimulus increases v_{post} , the distribution of the depolarization V changes in such a way that Q_b decreases and Q_a increases.

To illustrate the stochastic nature of the learning mechanism, we assume that the pre-synaptic spike train is Poisson distributed, while the afferent current to the post-synaptic neuron is a Gaussian distributed stochastic process. Such a situation is meant to mimic what happens during the presentation of a visual stimulus when the neuron is embedded in a large network (see Section 4).

The synaptic dynamics depends on the detailed time statistics of the spike trains of the pre- and post-synaptic neurons. During stimulation the synapses move tempo-

rarily up and down, triggered by the pre-synaptic spikes. Following the removal of the stimulus, the synaptic efficacy may return to its initial state, or it may make a transition to another state. Fig. 6 shows two cases at parity of mean spike rates of pre- and post-synaptic neurons for the synaptic dynamics that depends on the depolarization. In one case (left) a fluctuation drives the synaptic efficacy above threshold and, when the stimulus is removed, X is attracted to the high state: LTP has occurred. In the second case (right) when the stimulus is removed, X is below threshold and is attracted by the restoring current to the initial state. No transition occurred. In the two cases the statistics of the activity to be encoded—the mean spike frequency—is the same, but

the realization of the stochastic process that generated the pre- and the post-synaptic activity is different.

The synaptic dynamics is entirely deterministic and the stochasticity is provided by the neural activity. This mechanism should not be confused with other sources of randomness which might affect the long-term synaptic dynamics (and are not implemented here) or with the synaptic vesicle release of neurotransmitter which is an inherently stochastic process but it is not related to the dynamics of our internal variable X .

6.2.3. Dependence on spike-timing

Recent experiments [71] performed *in vitro* show that a synapse can be potentiated or depressed depending on the relative timing of the pre- and post-synaptic spikes. In particular, in the experiment [71] it was reported that if the pre-synaptic spike repeatedly precedes a post-synaptic action potential within a short time window (10–20 ms), the synapse is potentiated. If the opposite occurs (the post-synaptic action potential comes first), the synapse undergoes long-term depression. Such a mechanism for synaptic modifications is now usually termed spike timing dependent plasticity (STDP).

Many experiments have since studied STDP under various conditions as for the neural preparation, the protocol etc; the subject is still debated, and experimental evidence seems still not conclusive (see for example the review in [72]). In simulations, the authors of [73] found it useful to adopt a synaptic model that superficially resembles the above STDP, being explicitly dependent on spike timing. The motivation is not a claim of better biological plausibility, though we briefly speculate later about it, but computational convenience.

Fig. 8 sketches the specific synaptic dynamics we adopt for experimenting with dynamic learning scenar-

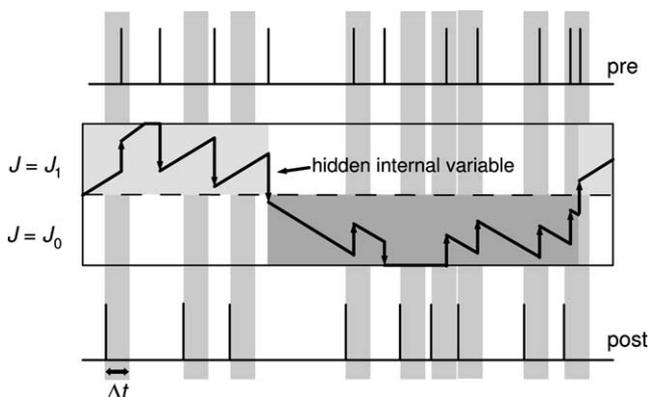


Fig. 8. An example dynamics of one of the spike-driven synapse used in the spiking neurons network to study the learning expression (adapted from [73]). The internal state $X(t)$ of the synapse is plotted as a function of time in the central plot. Above and below it are plotted the pre- and the post-synaptic spikes respectively. The synaptic threshold θ_x is represented by the dashed line. The shaded regions are the time intervals following post-synaptic spikes during which a pre-synaptic spike induces up-regulation of the internal synaptic variable X .

ios. In the figure we show a sample evolution of the analog internal synaptic variable (middle) for given pre-synaptic (top) and post-synaptic (bottom) spike trains, together with the associated time course of the synaptic efficacy (light gray strip: potentiated; dark gray strip: depressed). Temporary up (down) regulation, i.e. $X \rightarrow X + a$ ($X \rightarrow X - b$) of the internal variable occurs if a pre-synaptic spike comes within (beyond) Δt after the emission of a post-synaptic spike. A way to describe the mechanism is that each post-synaptic spike leaves an evanescent ‘trace’ for the purpose of the synaptic dynamics, and up-regulation occurs when the pre-synaptic spike finds the trace still there. A possible biophysical mechanism for such trace is described in [74].³

The synapse based on spikes timing undergoes stochastic transitions, analogously to what we have seen above for the synapse reading the post-synaptic depolarization: for given v_{pre} and v_{post} , the fluctuations in the inter-spike intervals can make a transition happen or not.

There is accumulating evidence [75] that a more realistic description of the protocols for inducing LTP and LTD probably requires a combination of the dependence on spike-timing—to take into account the effects of the back-propagating action potential—and of the dependence on the sub-threshold depolarization. A synaptic model which incorporates both dependences has been preliminarily studied in [70].

6.2.4. Encoding mean rates by reading the relative spike timing

While superficially reminiscent of the synaptic dynamics adopted in the spike-dependent synaptic plasticity (STDP) models inspired by the results in [71] and described above, the synapse adopted in [73] bears a major difference, related to the *variable size* of the time window for down-regulation of the synaptic variable. To get a feeling of why this is so, let us consider the statistics of synaptic up- and down-regulations for the present model, and a synapse subject to the same dynamics, but with *fixed* (though possibly different) time windows. Letting the pre-synaptic neuron fire at high rate, we examine the qualitative behavior of the fixed- and variable-window synapses, as a function of the post-synaptic emission rate v_{post} .

When v_{post} is low, the variable-window synapse gets almost always depressed, since almost always the ‘active traces’ of post-synaptic spikes have long expired; as v_{post}

³ Even though the synapse is introduced for purely computational motivations, as discussed below, it might be tempting to think of the trace as an expression of a back-propagated spike. On the other hand, to date it is not clear to what extent experimental data based on spike timings are able to constraint the allowed time window for LTD. The very definition of the time window without further specifications of the neural state seems to be a problematic concept (see [75]).

increases, probability of LTP increases at the expense of that for LTD, ending up with almost certain LTP for very high v_{post} . For the fixed-window synapse (at least for windows shorter than the average ISI), the scenario is totally different: for low v_{post} both the probabilities of LTP and LTD are low, and most of the times the synapse is in a ‘do-nothing’ condition. As v_{post} increases, both the probabilities for LTP and LTD increase, but keeping fixed their ratio, which only depend on the ratio between the windows. A Hebbian dependence on v_{post} is prevented in this case.

Two comments are in order. First, the above discussion applies if the trains of spikes of the pre- and post-synaptic neurons are statistically independent and the identity of the stimulus to be learnt is entirely encoded in the mean firing rates. Second, the discussion does not apply to other synaptic models with a suitable non-linear dependence on the rates, like the ones proposed in [76].

Fig. 9 illustrates the above discussion, reporting the curves of equal probability for LTP and LTD for a

single synapse with independent pre- and post-synaptic neurons emitting Poissonian spike trains. It is seen that for the synaptic models B (Section 6.2.1), C and D (Section 6.2.3) the $(v_{\text{pre}}, v_{\text{post}})$ plane is partitioned in well separated regions in which either LTP (high v_{pre} and v_{post}) or LTD (high v_{pre} and low v_{post}) dominate. The case of model A is shown to imply a superposition of LTP and LTD transitions, incompatible with a rate-dependent plasticity. Notice that as long as the time window available for down-regulation is variable (equal to the ISI minus the window for up-regulation) the qualitative behavior of such a synapse does not change much if the dependence on the relative spike timings is reversed.

While the interpretation of the available data on long-term synaptic modifications is still controversial, the theoretical discussion in Section 5.3 corroborates the long standing belief that a Hebbian-like learning term would be required for building up associative memory in a quite general scenario and there are reasons to keep, at least as a useful working hypothesis, the ‘Hebbian’ rate dependence of long-term synaptic modifications.

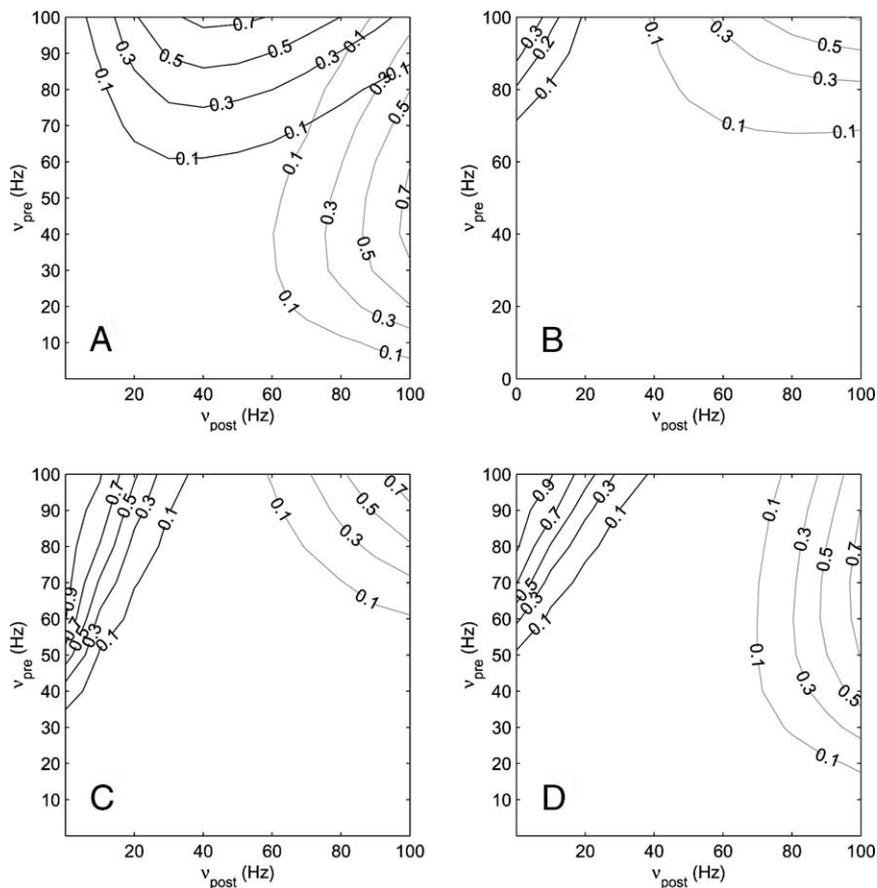


Fig. 9. Curves of iso-probability of potentiation (gray) and depression (black) for synapses following the rule suggested by the experiment by Markram et al. (A), the model of Fusi et al. (B), the synapse with a variable window for depression used in [73] (C) and a synapse with a variable window for depression, but otherwise the same as A (D), driven by pre-synaptic and post-synaptic neurons emitting Poisson spike trains respectively at frequencies v_{pre} and v_{post} .

7. General consideration on the coupled dynamics of neurons and synapses

We now take a step towards the study of the synaptic dynamics as it autonomously emerges from the ongoing neural activity in the network, facing the full complexity of the mutual dependence between the neural and synaptic degrees of freedom.

Even at the first, superficial level of analysis, the inclusion of synaptic dynamics has profound implications for the collective behavior of the network. First, neural and synaptic dynamics are coupled in a very intricate way: synaptic modifications are driven by neural activities, which are in turn determined by the pattern of synaptic connections at any given time; the two sets of dynamical variables in this loop are generally assumed to be associated with very different characteristic time scales (*fast* neural variables, as opposed to the *slow* synaptic ones). ‘Learning’ and ‘retrieval’ are no longer separate categories: an incoming stimulus will act both as a trigger for the retrieval mechanism (which will be associated with relaxation into the attractor providing the internal representation of that stimulus), and a boost for synaptic modification that shapes “tuning curves” and forms selective delay activity; the balance between the two actions will depend in general on many factors, including the familiar or not-familiar nature of the stimulus, and the degree of overlap among the stimuli embedded as memories in the synaptic matrix.

The very fact that the synaptic state of the network depends on the neural activities raises new questions about the stability of any neural state: it is easy to imagine situations in which the synaptic matrix keeps changing (perhaps at different rates, depending on the various neural regimes of spontaneous activity, stimulation or collective, selective reverberation), causing a drift in the global state of the network, in which all memories can be doomed to fade away, or local neural activities can ignite explosions of activity throughout the network, or any other kind of behavior can appear, that makes any computationally relevant collective state unstable.

Alternatively, short- and long-term transients on the synaptic values may produce interesting temporal features in the neuronal dynamics.

7.1. Constraints related to Hebbian learning

The typical learning scenario has the network initially ‘unstructured’: all neurons in the network share the same statistical distribution of synaptic efficacies on their dendritic trees. Synaptic modifications are assumed to take place only during stimulation. Indeed, if the synaptic structure has to embody *memory*, such as to code for the familiarity of an incoming stimulus on the basis of its past experience, it is natural to assume that

changes are only brought about by stimuli. The structured network should be able to keep its memory store essentially unchanged for long times, in the absence of incoming stimuli; if spontaneous activity could modify the synapses, or the synaptic matrix were to change during the reverberating, attractor state expressing its recognition of a familiar stimulus, either only the memory of the last stimulus would survive, or the memory would be wiped out altogether. Furthermore only those stimuli which consistently characterize the network interaction with its environment should have a chance to significantly affect the synaptic structure, while those occasionally occurring should appear as noise in the process of synaptic structuring: in most cases, for the building up of working memory states the *learning rate* has to be small and fairly constant in time, as we discussed in Section 5.2.

From the above observations general constraints follow, to be satisfied by the coupled dynamics of neurons and synapses in a dynamic learning scenario. Let v_{spon} , v_{stim} and v_{sel} denote respectively the average spike emission rate of neurons in the spontaneous state, of the ‘visual response’ of neurons directly affected by a stimulation, and of neurons that, as a result of a process of stimulus-related synaptic structuring, have a sustained, stimulus-specific activity after removal of the stimulus. In a Hebbian scenario of rate-dependent synaptic potentiation (we will discuss later LTD), for synaptic modifications to take place starting from an unstructured network v_{stim} , much greater than v_{spon} , has to be high enough to induce potentiation in the synapses connecting stimulated neurons, but still low enough to have slow learning. In the initial stage $v_{\text{sel}} = v_{\text{spon}}$. As a subset of synapses start getting potentiated, in general v_{stim} increases: indeed, the average afferent current of the stimulated neurons will ultimately have to increase enough to support the reverberant state. But if v_{stim} grows too fast, the desired feature of slow learning is lost, and the learning rate start increasing, making recent stimuli more and more effective in leaving their trace in the synaptic structure, at the expense of previous stimuli. Then, when the synaptic structure becomes able to support stimulus-specific reverberant states, a further constraint appears on v_{sel} , for if v_{sel} were greater than v_{stim} at the beginning of learning, this would imply an ongoing synaptic modification during the reverberation: the selective delay activity would tend to reinforce the memory trace of the last stimulus far beyond the desired extent.

Thus, summarizing the above observations, Hebbian dynamic learning has to cope with subtle problems of stability of the network states, and stability of the learning process.

Recently there have been first attempts to implement the learning dynamics in large scale simulations of integrate-and-fire neurons and plastic, spike-driven synapses of the type described in Section 6 [69,73].

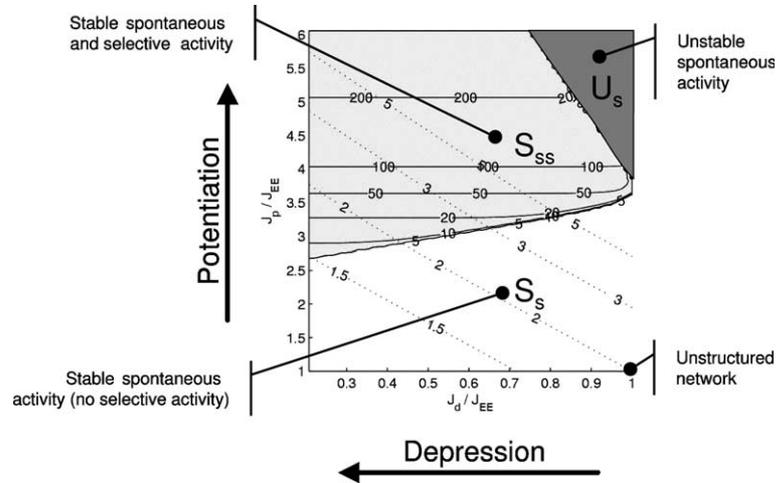


Fig. 10. The potentiation-depression (PD) plane which represents the stable asynchronous states at different levels of potentiation and depression of synaptic efficacies due to a learning process involving a finite set of uncorrelated stimuli (see text for details, adapted from [73]).

In the following we provide an expanded account of the findings in Ref. [73] where the authors investigate the formation of working memory with a network of linear integrate-and-fire neurons (i.e. with constant leakage and a rigid lower boundary) and the dynamic spike-driven synapses introduced in Section 6.2.3.

7.2. Exploratory tools for dynamic learning

From the discussion above it follows that scenarios for dynamic learning would better imply a sequence of quasi-equilibrium states for the coupled system of neurons and synapses, to endow the system with appropriate stability properties. In such a scenario, the equilibrium mean field theory for the neural network with fixed synapses would hopefully provide a good first approximation in the proximity of each (stable) point along the learning trajectory, and would therefore serve as a useful exploratory tool to spot interesting regions in the system's state space to look into.

For a network with N neurons the analysis of the coupled neural and synaptic dynamics requires to solve a system of $O(N + N^2)$ coupled equations. Standard approaches to the numerical integration of those equations imply a prohibitive $O(N^3)$ computational load. The more so, if one considers that the characteristic times of the synaptic dynamics are inherently slow, and one is faced with the problem of accurately reproduce an interacting system with one *fast* (neural) and one *slow* (synaptic) sets of dynamic variable.

We therefore use the event-driven approach to simulations introduced in [77], which lowers the complexity to $O(N^2)$.⁴

Before embarking in large scale simulations, however, some indications should be drawn from the theory. The coupled neural and synaptic system has in fact very many parameters, and a blind exploration of the system's state space is in most cases hopeless.

Even if no real 'double dynamics' theory is available, in the regime of slow learning of interest a 'good learning trajectory' of the network (i.e. the sequence of synaptic configurations travelled by the network on its way to the onset of selective reverberant activity) should be a sequence of quasi-equilibrium states, meaning that the rate of synaptic changes should be low enough to make the usual mean field predictions for fixed synapses approximately valid for the real dynamic network happening to have the same average synaptic efficacies at a given stage of its learning history.

In order to prepare for an informed exploration of the system state space in simulation, one can then characterize the 'quasi-static' system by suitably parameterizing the synaptic configurations, and for each set of parameters predicting the available stable fixed points according to the "extended mean field theory" approach [5].

Assuming that synapses undergoing Hebbian potentiation and those homosynaptically depressed all share the same average efficacies (J_p and J_d respectively), a compact parameterization (extending the one adopted in [5]) can be devised, exposing in the plane (J_p, J_d) the nature of the stable collective states predicted by MFT, as shown in Fig. 10.

For each point in the plane, MFT predicts the allowed states and assesses their stability, thus providing, for example, an indication as to the region of parameters where the onset of the simultaneously stable spontaneous and selective states, or the one for which the spontaneous state gets destabilized because of excessive potentiation.

In Fig. 10 J_{EE} is the average efficacy of the unaffected synapses. Dotted (solid) contour lines are iso-frequency

⁴ The method of [77] has been improved in [78].

curves for spontaneous (selective) activity. For white regions (marked S_S) MFT predicts only stable spontaneous states; for light gray regions (marked S_{SS}) there are both spontaneous and selective stable states, while for the dark gray region (marked U_S) the spontaneous activity is unstable.

Successful learning trajectories will bring the network from the totally unstructured synaptic state, lower right corner in the plane, to a state supporting a low rate, unspecific “spontaneous” state coexisting with a selective, higher rate, state to be interpreted as the neural substrate of a working memory state (light gray region in the plane).

8. Phenomenology of dynamic learning

Fig. 11 illustrates a successful learning history. The actual trajectory of the system as it is generated by the simulation is superimposed on the mean field predictions in the PD-plane. Each diamond is one of the synaptic states during the learning history.

The stimulation protocol is as follows: non-overlapping sets of neurons are stimulated (their external afferents increase their emission rate) by five stimuli, which are presented in sequence. At the beginning, starting from a totally unstructured synaptic distribution, after removal of each stimulus the neural activity relaxes back to the spontaneous state.

Gradually, LTP and LTD start driving the network towards the desired gray region where spontaneous and selective states coexist. The upper snapshot illustrates the scenario in the mature network: after removal of each stimulus a neural population largely overlapping the one involved in the stimulation relaxes to the reverberant selective state. Notice how in this case the simulation matches well the predictions from mean field theory, all along the learning trajectory, supporting the assumption that for slow enough learning the network lives in quasi-equilibrium states.

In Fig. 12 is illustrated a schematic plot of how the synaptic matrix is modified during the learning stage. The evolution of J_p and J_d can be described by a non-stationary Markov process. From the figure is clear how the learning process is self-regulated: After a high enough number of stimuli presentations the LTP and LTD levels have not further change from the asymptotic values.

8.1. LTP-LTD balance

Let us take the point of view of a neuron (one of those involved in learning, i.e. one neuron affected by one stimulus)⁵ looking at the changes in the synaptic

⁵ Each neuron can be stimulated by just one stimulus, because of the simplifying choice we made to have ‘orthogonal’ stimuli.

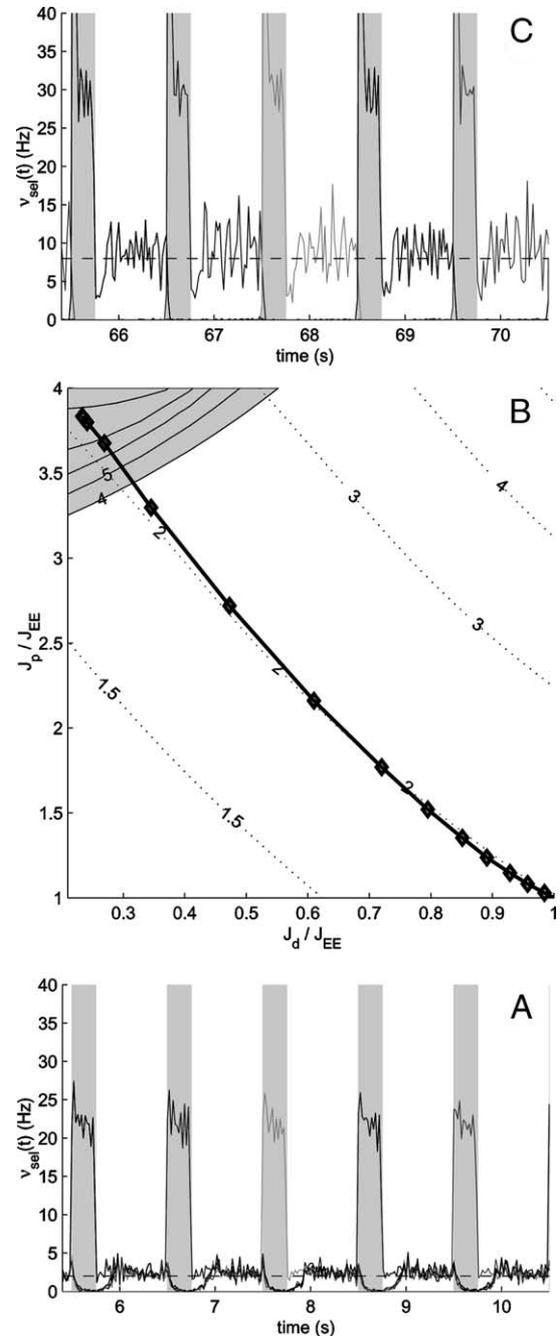


Fig. 11. A successful learning trajectory. (A) Neural activity of the five uncorrelated stimuli used in the learning process for a totally unstructured synaptic distribution (starting point for the simulation): No selective delay activity shown after visual response. (B) The PD-plane for the simulated network superimposed to the learning trajectory emerging from the simulation: Each diamond indicates the system state after a presentation of the entire set of stimuli. (C) As in (A), at the end of the simulation the system reaches the region of the PD-plane where spontaneous and selective states coexist: the selective delay activity following a stimulation is clearly visible. See text for details (this figure is adapted from [73]).

connections on its dendritic tree. At the beginning of learning the synaptic contacts are homogeneously dis-

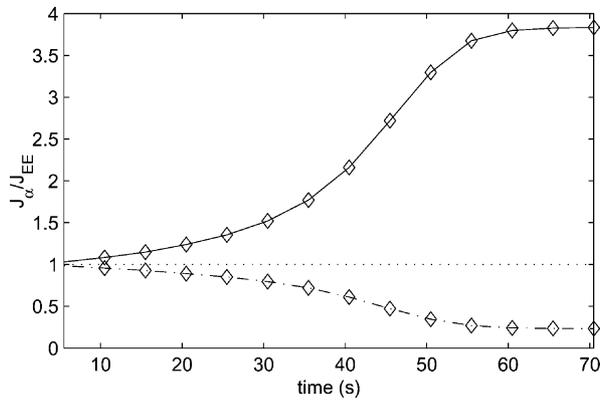


Fig. 12. Potentiation level J_p/J_{EE} (solid curve) and depression level J_d/J_{EE} (point-dashed curve) during a learning session. Data are from the same simulation shown in Fig. 11.

tributed. In the desired scenario, as the stimulus driving that neuron start being presented, the dendritic tree under consideration start partitioning into three regions: a region of essentially unaffected synapses, whose pre-synaptic neurons belong to the ‘background’ population of neurons which are never stimulated (their emission rate is small, so they would not trigger synaptic modifications); a region of depressed synapses, whose pre-synaptic neurons are stimulated by other stimuli (in this case the homosynaptic LTD mechanism is at work, since when one of the stimuli non-affecting the considered neuron is active, its pre-synaptic stimulated neurons find the neuron essentially inactive, and the corresponding synapse gets depressed); a region of potentiated synapses, whose pre-synaptic neurons are stimulated by the same stimulus as the neuron under consideration.

It is relevant to look at the average changes in the afferent current to the neuron under consideration in the above conditions. Clearly, the ‘unaffected’ dendritic tract does not bear changes, while the other two regions bring about conflicting effects (the depressed one lowering the average afferent current, the potentiated one increasing it). It is precisely the balance of those conflicting pushes which is necessary in order to allow for simultaneously stable spontaneous and stimulus selective states. For example, if the increase in the afferent current due to the potentiated dendritic tract is not compensated by enough LTD in the depressed tract, the network runs the risk of destabilizing the spontaneous state; in the opposite case, the selective reverberant state might not be produced because the insufficient increase in the potentiated dendritic tract prevents supporting it.

Hence the parameters of the synaptic model must be tuned to guarantee a balance between LTD and LTP at each learning stage (see e.g. Fig. 12). The introduction of a suitable regulatory mechanism would probably relax this strict requirement (see Section 8.3).

The balance between LTP and LTD guarantees also that more complex, overlapping random patterns, are

learnt and forgotten in an optimal way. The patterns can be learnt at the maximum speed compatible with a memory span which comprises all the patterns that should be stored. Indeed, the memory span becomes optimal when the palimpsest term (see Eq. (6)) tends to 1, i.e. when for instance the LTP transition probability multiplied by the probability of having a pair of simultaneously active neurons (the condition for inducing LTP) is as small as the LTD probability multiplied by probability that the pre-synaptic neuron is active and post-synaptic neuron is inactive. In the case of random patterns with a small fraction f of active neurons the LTD probability should be f times smaller than the LTP probability to ensure the balance [56].

8.2. Finite-size effects

In a network with N neurons, if $N \rightarrow \infty$ all neurons in a given population (i.e. those sharing the same statistical properties of the afferent current, for example neurons affected by the same stimulus) fire at the same rate, with small fluctuation from neuron to neuron.

For finite N finite-size effects are important, and bring about deviations from the predictions of the mean field theory. In particular, for finite N each population have a *distribution* of emission rates. Let us consider the synapses connecting populations of neurons stimulated by the same stimulus, and therefore supposed to get potentiated: the high- and low-rate tails of the actual frequency distribution corrupt the homogeneity of the pattern of synaptic transition probabilities, such that in the same synaptic group some synapses will have too high LTP probability, while others will be unexpectedly unchanged. Similarly, finite-size effects can provoke unwanted synaptic transitions where they are not expected and harmful (such as a potentiation of synapses involving post-synaptic background neurons, which can become the seed of instability for the spontaneous state). One ingredient which makes finite-size effects more or less harmful is the character of the ‘synaptic transfer function’, meaning the function giving the LTP/LTD transition probabilities as functions of the pre- and post-synaptic emission rates. The sensitivity of this function in the critical region where the rate distributions involved overlap is an important factor in determining how serious finite-size effects are going to be.

Fig. 13 provides an illustration. The plots report the fraction of potentiated synapses on the dendritic tract of a stimulated neuron the synapses of which are supposed to potentiate. The left panel, top to bottom, shows a ‘bad’ situation, in which finite-size effects split the synaptic distribution and end up with a bi-modal distribution, implying a significant deviation from mean field predictions, and possibly hindering the formation of the reverberant state. The right panel shows how a change

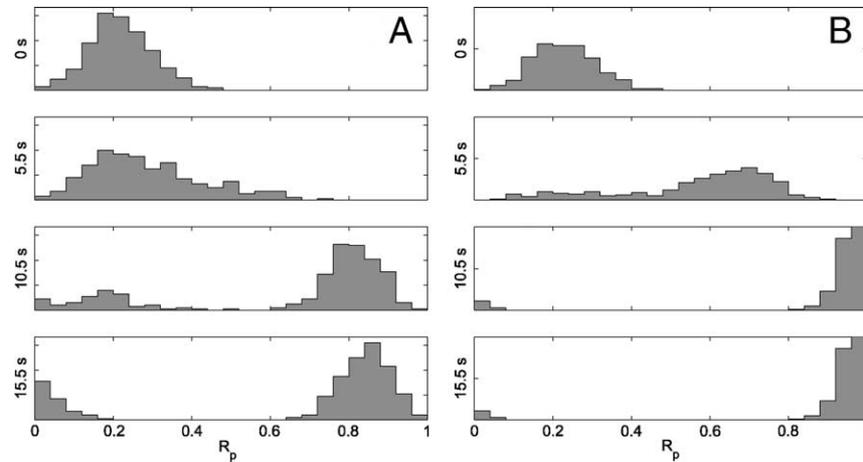


Fig. 13. Histograms of the fraction R_p of potentiated synapses a neuron receives from homogeneous cells in a stimulated sub-population at different successive learning stages. A learning stage consists of a sequence of five different stimuli, each presented for 250 ms every second. Each histogram is sampled at the times reported on the left of the histograms. (A) Case in which the dendritic tree develops a 'bad' structure (some neurons express LTD rather than LTP); (B) A well separated case.

in the synaptic parameters (for the same network) allows the desired synaptic structuring.

It is because of its lesser sensitivity to finite-size effects that we adopted in most simulation, the synaptic model described in Section 6.2.3; it should be kept in mind, however, that synapse and the one based on the post-synaptic depolarization produce the same learning scenario, being both meant to implement the same dependence of LTP and LTD on the pre- and post-synaptic rates [69].

8.3. Regulatory mechanisms

Regulatory mechanisms that prevent the excitatory activity from growing boundedly turned out to be essential to guarantee the stability of the network dynamics throughout the learning process. One inherent limiting factor for the growth of the mean spike frequencies is synaptic saturation: the internal synaptic state and, hence, the synaptic efficacies are limited from below and from above. When all synapses saturate, e.g. in the potentiated state, the synaptic structure cannot change any longer in the case of non-overlapping patterns. This mechanism prevents the frequencies of the stimulated neurons from growing any further and the learning process actually stops. However, even in this case an extra regulatory mechanism is required to guarantee that there exists a trajectory of stable points that connects the initial state to the final, saturated one. Among all the other important considerations that have been discussed, one seems to be essential in both studies [69,73]: the mean frequency under stimulation should be kept almost constant throughout the learning process. This requires a regulatory mechanism that keeps under control the increase of the mean spike frequencies due to the growth of the synaptic feedback. In [69] the authors

tune manually the external afferent frequencies in order to keep the response under stimulation almost constant. In this context, homeostatic mechanisms like those reviewed in [79] might play a role. They essentially decrease the external input every few stimuli presentations in order to compensate for the increased synaptic recurrent feedback. In [73] the authors implement an automatic mechanism based on short-term depression (STD) [80,81] for decreasing the effect of recurrent synapses connecting stimulated neurons (J_p) when the spike frequency under stimulation tends to grow too quickly. Each synaptic efficacy decreases temporarily upon the arrival of a pre-synaptic spike. It then recovers with some typical time constant of the order of tens or hundreds of milliseconds. If other spikes arrive from the pre-synaptic side during the recovery phase, the synaptic efficacy is depressed even more and eventually might become totally ineffective. This mechanism is meant to mimic the average effect of the temporary changes of the synaptic neurotransmitter release probability. In Fig. 14 we compare the frequency of the stimulated population in the presence and in the absence of STD. If the increase of J_p is not adequately balanced by increased depression, the spontaneous activity is easily destabilized, and the equilibrium can be very critical.

Possible alternative mechanisms might involve neural adaptation or modifications of the synaptic dynamics. In the first case a negative feedback limits the output frequency of the pre-synaptic neurons and the regulatory mechanism is probably equivalent to the short depression described above. In the second case one should introduce a new element in the synaptic dynamics that blocks learning when the frequency of the post-synaptic neuron becomes too high. Such a mechanism is currently investigated and some preliminary analysis shows that it would also have the additional

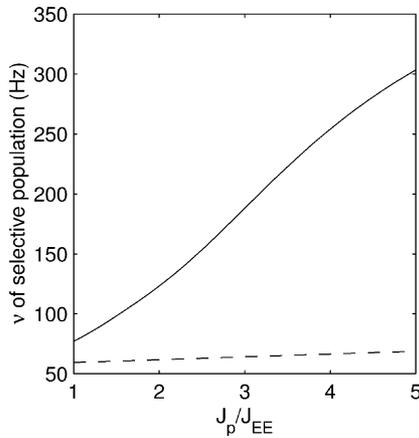


Fig. 14. An example of visual response during a learning history with (dashed line) and without (solid line) synaptic STD.

advantage of permitting the storage of more complex, correlated patterns of activity [70]. In the end the regulatory system in the brain is probably a mixture of all these mechanisms.

9. Discussion

9.1. Statistics of the stimuli

In a more realistic case, less constrained statistics of stimuli (overlapping random stimuli, stimuli spanning different coding levels, correlated stimuli, complex protocols of presentation) should be considered. The ‘feasibility proofs’ described in [69,73], though instructive, refer to very unrealistic scenarios as for the spatial and temporal statistics of the incoming stimuli: two different stimuli elicit responses in disjoint sets of neurons—no overlap between stimuli; all the stimuli share the same *coding level* (the fraction of neurons excited by the stimulus); stimuli are presented in very tame temporal sequences: identical repetitions of the same sequence of stimuli. A non-zero overlap between stimuli makes the learning process very different: the generic synapse feels, during learning, conflicting pushes towards potentiation or depression, depending on the current stimulus and the learning process can be described as a random walk between the two stable states when the stimuli are random (see Section 5.3). The synaptic configuration never comes to rest and, even in case of successful learning, keep on fluctuating around a non-saturated configuration with potentiated and depressed synapses coexisting in the same synaptic set. In case of random patterns with the same average coding level, the stochastic selection described in Section 5.2 allows the network to perform well as an associative memory and the storage capacity can be optimal provided that learning is slow enough. Even if not needed, in the case of non-overlapping

stimuli studied in [69,73] and reviewed here the stochastic selection mechanism is already implemented: each stimulus induces long-term modifications (the internal synaptic variable crosses the threshold) in a randomly selected subset of synapses. The development of the distribution of the synapses can be approximately described as the Markov process described in Section 5.3. The analysis of the network dynamics for overlapping random stimuli is much more complicated and probably will require extra regulatory mechanisms to guarantee the stability also in the presence of the interference between different stimuli. However all the basic elements for storing an extensive number of random patterns are already contained in the analysis summarized here.

Relaxing the constraint of equal coding level for all patterns might bring about a complex learning history, in which the speed of learning (and forgetting) is very different for different patterns. Also, allowing for arbitrarily complex temporal sequences of stimuli makes the learning process more complex, but would allow interesting comparisons with experimentally tested protocols.

9.2. Lifetime of attractors (multi-stability)

Since the number N of neurons in the network is finite, the attractor states are meta-stable: they can persist for long time, but the finite-size fluctuations, even in the absence of external perturbations, will ultimately drive the network outside the basin of attraction. In other words, the barriers which separate the attractor fixed points are finite for finite N , and whatever consequent finite amount of noise will make the system cross them sooner or later. But knowing if it is typically sooner or later in situations of interest is important. First of all, in ‘realistic’ conditions the network lives in a noisy environment, e.g. we should envisage the possibility that while the attractor reverberation is active *distractor* stimuli intervene, possibly unfamiliar or weak but potentially able to make the attractor state decay (see [41] for an example). Besides, when the attractor decays because of intervening stimuli, it is relevant to know how quickly this happens (see next item). Another example is the observed spontaneous alternation of attention in psychophysics experiments, with random distribution of the times of shift: the so-called multi-stable perception.

9.3. Miyashita correlations: strong prediction of the model, hard to implement in realistic modelling

Experiments show [11,27] that in a DMS task, with visual stimuli learnt in a fixed order, the neural representations (selective sustained delay activity states) of stimuli which are neighbors in the training sequence are

correlated. In other words a conversion from *temporal* to *spatial* correlation occurs [82,83]. In the attractor picture, one interpretation is that for two successive stimuli, upon presenting the second the reverberation induced by the first is still active, and synaptic potentiation can take place, for synapses connecting neurons participating in the dying reverberation, to neurons starting to visually respond to the second stimulus. In this scenario the characteristic times of the decay of the old attractor, and those of the onset of the new visual response, are crucial ingredients for modelling the phenomenon.

References

- [1] D.J. Amit, The Hebbian paradigm reintegrated: local reverberations as internal representations, *Behav. Brain Sci.* 18 (1995) 617–657.
- [2] D.O. Hebb, *The Organization of Behavior: A Neuropsychological Theory*, John Wiley & Sons, 1949.
- [3] T.J. Sejnowski, Storing covariance with nonlinearly interacting neurons, *J. Math. Biol.* 4 (1977) 303.
- [4] J.J. Hopfield, Neural networks and physical systems with emergent selective computational abilities, *Proc. Natl. Acad. Sci. USA* 79 (1982) 2554.
- [5] D.J. Amit, N. Brunel, Model of global spontaneous activity and local structured (learned) delay activity during delay periods in cerebral cortex, *Cereb. Cortex* 7 (1997) 237–252.
- [6] D.J. Amit, N. Brunel, Dynamics of a recurrent network of spiking neurons before and following learning, *Network* 8 (1997) 373–404.
- [7] J. Fuster, Unit activity in prefrontal cortex during delayed-response performance: neural correlates of transient memory, *J. Neurophysiol.* 36 (1973) 61.
- [8] H. Niki, Prefrontal unit activity during delay alternation in the monkey, *Brain Res.* 68 (1974) 185.
- [9] J.M. Fuster, Inferotemporal units in selective visual attention and short-term memory, *J. Neurophysiol.* 64 (1990) 681–697.
- [10] J.M. Fuster, J.P. Jervey, Inferotemporal neurons distinguish and retain behaviorally relevant features of visual stimuli, *Science* 212 (1981) 952–955.
- [11] Y. Miyashita, H.S. Chang, Neuronal correlate of pictorial short-term memory in the primate temporal cortex, *Nature* 331 (1988) 68.
- [12] Y. Miyashita, Neuronal correlate of visual associative long-term memory in the primate temporal cortex, *Nature* 335 (1988) 817.
- [13] K. Sakai, Y. Miyashita, Neural organization for the long-term memory of paired associates, *Nature* 354 (1991) 152–155.
- [14] E.K. Miller, L. Li, R. Desimone, Activity of neurons in anterior inferior temporal cortex during a short-term memory task, *J. Neurosci.* 13 (1993) 1460–1478.
- [15] K. Nakamura, K. Kubota, Mnemonic firing of neurons in the monkey temporal pole during a visual recognition memory task, *J. Neurophysiol.* 74 (1995) 162–178.
- [16] S. Funahashi, C.J. Bruce, P.S. Goldman-Rakic, Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex, *J. Neurophysiol.* 61 (1989) 331.
- [17] F.A.W. Wilson, S.P.O. Scalaidhe, P.S. Goldman-Rakic, Dissociation of object and spatial processing domains in primate prefrontal cortex, *Science* 260 (1993) 1955–1958.
- [18] E.K. Miller, C. Erickson, R. Desimone, Neural mechanisms of working memory in prefrontal cortex of the macaque, *J. Neurosci.* 16 (1996) 5154–5167.
- [19] S.P.O. Scalaidhe, F.A.W. Wilson, P.S. Goldman-Rakic, Areal segregation of face-processing neurons in prefrontal cortex, *Science* 278 (1997) 1135–1138.
- [20] S.C. Rao, G. Rainer, E.K. Miller, Integration of what and where in the primate prefrontal cortex, *Science* 276 (1997) 821–824.
- [21] K.W. Koch, J.M. Fuster, Unit activity in monkey parietal cortex related to haptic perception and temporary memory, *Exp. Brain Res.* 442 (1989) 359.
- [22] J.M. Fuster, *Memory in the Cerebral Cortex: An Empirical Approach to Neural Networks in the Human and Nonhuman Primate*, MIT Press, 1995.
- [23] Y. Miyashita, T. Hayashi, Neural representation of visual objects: encoding and top-down activation, *Curr. Opin. Neurobiol.* 10 (2000) 187–194.
- [24] X.J. Wang, Synaptic reverberation underlying mnemonic persistent activity, *Trends Neurosci.* 24 (2001) 455–463.
- [25] E.A. Murray, D. Gaffan, M. Mishkin, Neural substrates of visual stimulus–stimulus association in rhesus monkeys, *J. Neurosci.* 13 (1993) 4549–4561.
- [26] S.A. Gutnikov, M. Yuan-Ye, D. Gaffan, Temporo-frontal disconnection impairs visual–visual paired association learning but not configural learning in macaca monkeys, *Euro. J. Neurosci.* 9 (1997) 1524–1529.
- [27] V. Yakovlev, S. Fusi, E. Berman, E. Zohari, Inter-trial neuronal activity in inferior temporal cortex: a putative vehicle to generate long-term visual associations, *Nature Neurosci.* 1 (4) (1998) 310–317.
- [28] R. Romo, C.D. Brody, A. Hernandez, L. Lemus, Neuronal correlates of parametric working memory in the prefrontal cortex, *Nature* 399 (1999) 470–473.
- [29] H.S. Seung, How the brain keeps the eyes still, *Proc. Natl. Acad. Sci. USA* 93 (1996) 13339–13344.
- [30] K. Tanaka, Inferotemporal cortex ABD object vision, *Ann. Rev. Neurosci.* 19 (1996) 109–140.
- [31] D.J. Amit, S. Fusi, V. Yakovlev, Paradigmatic attractor cell, *Neural Comput.* 9 (1997) 1071–1093.
- [32] Y. Miyashita, Inferior temporal cortex: where visual perception meets memory, *Ann. Rev. Neurosci.* 16 (1993) 245–263.
- [33] E. Kobatake, G. Wang, K. Tanaka, Effects of shape-discrimination training on the selectivity of inferotemporal cells in adult monkeys, *J. Neurophysiol.* 80 (1998) 324–330.
- [34] C.A. Erickson, R. Desimone, Responses of macaque perirhinal neurons during and after visual stimulus association learning, *J. Neurosci.* 19 (1999) 10404–10416.
- [35] H. Tomita, M. Ohbayashi, K. Nakahara, I. Hasegawa, Y. Miyashita, Top down signal from prefrontal cortex in executive control of memory retrieval, *Nature* 401 (1999) 699–703.
- [36] D.J. Amit, M. Tsodyks, Quantitative study of attractor neural network retrieving at low spike rates: I. Substrate-spikes, and neuronal gain, *Network* 2 (1991) 259.
- [37] D.J. Amit, M. Tsodyks, Effective and attractor neural networks in cortical environment, *Network* 3 (1992) 121–137.
- [38] P. Dayan, L.F. Abbott, *Theoretical Neuroscience*, MIT Press, 2001.
- [39] W. Gerstner, W. Kistler, *Spiking Neuron Models*, Cambridge University Press, Cambridge, 2002.
- [40] S. Fusi, M. Mattia, Collective behavior of networks with linear (VLSI) integrate-and-fire neurons, *Neural Comput.* 11 (3) (1999) 633–653.
- [41] N. Brunel, X.J. Wang, Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition, *J. Comput. Neurosci.* 11 (1) (2001) 63–85.
- [42] L.M. Ricciardi, L. Sacerdote, The Ornstein–Uhlenbeck process as a model for neuronal activity. I. Mean and variance of the firing time, *Biol. Cybern.* 35 (1) (1979) 1–9.
- [43] N. Brunel, Persistent activity and the single cell *f–i* curve in a cortical network model, *Network* 11 (2000) 261–280.

- [44] A. Rauch, G. La Camera, H.-R. Luescher, W. Senn, S. Fusi, Neocortical pyramidal cells respond as integrate-and-fire neurons to in vivo-like input currents, *J. Neurophys.* 90 (2003) 1598–1612.
- [45] A. Treves, Mean-field analysis of neuronal spike dynamics, *Network* 4 (1993) 259–284.
- [46] L.F. Abbott, C. van Vreeswijk, Asynchronous states in networks of pulse-coupled oscillators, *Phys. Rev. E* 48 (2) (1993) 1483–1490.
- [47] W. Gerstner, Time structure of the activity in neural network models, *Phys. Rev. E* 51 (1) (1995) 738–758.
- [48] N. Brunel, V. Hakim, Fast global oscillations in networks of integrate-and-fire neurons with low firing rates, *Neural Comput.* 11 (1999) 1621–1671.
- [49] N. Brunel, Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons, *J. Comput. Neurosci.* 8 (2000) 183–208.
- [50] M. Mattia, P. Del Giudice, Population dynamics of interacting spiking neurons, *Phys. Rev. E* 66 (2002) 051917.
- [51] M. Mascarò, D.J. Amit, Effective neural response function for collective population states, *Network* 10 (1999) 351–373.
- [52] S. Fusi, Hebbian spike-driven synaptic plasticity for learning patterns of asynchronous activity, *Biol. Cybern.* 87 (5–6) (2002) 459–470.
- [53] G. Parisi, A memory which forgets, *J. Phys. A: Math. Gen.* 19 (1986) L617.
- [54] J.P. Nadal, G. Toulouse, J.P. Changeux, S. Dehaene, Networks of formal neurons and memory palimpsests, *Europhys. Lett.* 1 (1986) 535.
- [55] D.J. Amit, S. Fusi, Constraints on learning in dynamic synapses, *Network* 3 (1992) 443.
- [56] D.J. Amit, S. Fusi, Learning in neural networks with material synapses, *Neural Comput.* 6 (5) (1994) 957–982.
- [57] D.J. Amit, *Modeling Brain Function*, Cambridge University Press, NY, 1989.
- [58] F.P. Battaglia, S. Fusi, Learning in neural networks with partially structured synaptic transitions, *Network* 6 (1995) 261.
- [59] P.E. Kahn, K.Y.M. Wong, D. Sherrington, A memory model with novel behaviour in sequential learning, *Network* 6 (1995) 415–427.
- [60] N. Brunel, F. Carusi, S. Fusi, Slow stochastic Hebbian learning of classes of stimuli in a recurrent neural network, *Network* 9 (1) (1998) 123–152.
- [61] S. Fusi, Prototype extraction in material attractor neural networks with stochastic dynamic learning, In: S.R. Orlando, D. Ruck (Eds.), *Proceedings of SPIE 95*, vol. II, 1995, pp. 1027–1038.
- [62] W.R. Softky, C. Koch, The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs, *J. Neurosci.* 1 (1993) 334–350.
- [63] S. Fusi, M. Annunziato, D. Badoni, A. Salamon, D.J. Amit, Spike-driven synaptic plasticity: theory, simulation, VLSI implementation, *Neural Comput.* 12 (2000) 2227–2258.
- [64] E. Chicca, S. Fusi, Stochastic synaptic plasticity in deterministic a VLSI networks of spiking neurons, in: F. Rattay (Ed.), *Proceedings of the World Congress on Neuroinformatics, ARGESIM/ASIM Verlag, Vienna, 2001*, pp. 468–477.
- [65] M. D’Andreagiovanni, V. Dante, P.D. Giudice, M. Mattia, G. Salina, Emergent asynchronous, irregular firing in a deterministic analog VLSI recurrent network, in: F. Rattay (Ed.), *Proceedings of the World Congress on Neuroinformatics, ARGESIM/ASIM Verlag, Vienna, 2001*.
- [66] S. Fusi, P. Del Giudice, D.J. Amit, Neurophysiology of a VLSI spiking neural network: Lann21, in: S.-I. Amari, L. Giles, M. Gori, V. Piuri (Eds.), *Proceedings of the IEEE-INNS-ENNS International Joint Conference on Neural Networks*, vol. III, 2000, pp. 121–126.
- [67] T.V.P. Bliss, G.L. Collingridge, A synaptic model of memory: long term potentiation in the hippocampus, *Nature* 361 (1993) 31–39.
- [68] D.V. Buonomano, M.M. Merzenich, Cortical plasticity: from synapses to maps, *Ann. Rev. Neurosci.* 21 (1998) 149–186.
- [69] D.J. Amit, G. Mongillo, Spike-driven synaptic dynamics generating working memory states, *Neural Comput.* 15 (2003) 565–596.
- [70] S. Fusi, Spike-driven synaptic plasticity for learning correlated patterns of mean firing rates, *Rev. Neurosci.* 14 (2003) 73–84.
- [71] H. Markram, J. Lubke, M. Frotscher, B. Sakmann, Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs, *Science* 275 (1997) 213–215.
- [72] L.F. Abbott, S.B. Nelson, Synaptic plasticity: taming the beast, *Nature Neurosci.* 3 (2000) 1178–1183.
- [73] P. Del Giudice, M. Mattia, Long and short-term synaptic plasticity and the formation of working memory: a case study, in: J. Bower (Ed.), *Neurocomputing—Computational Neuroscience: Trends in Research 2001*, vols. 38–40, Elsevier Science, Amsterdam, 2001, pp. 1175–1180.
- [74] W. Senn, H. Markram, M. Tsodyks, An algorithm for modifying neurotransmitter release probability based on pre- and postsynaptic spike timing, *Neural Comput.* 13 (1) (2001) 35–67.
- [75] P.J. Sjöström, G.G. Turrigiano, S.B. Nelson, Rate, timing and cooperativity jointly determine cortical synaptic plasticity, *Neuron* 32 (2001) 1149–1164.
- [76] W. Senn, Beyond spike-timing: the role of nonlinear plasticity and unreliable synapses, *Biol. Cybern.* 87 (5–6) (2002) 344–355.
- [77] M. Mattia, P. Del Giudice, Efficient event-driven simulation of large networks of spiking neurons and dynamical synapses, *Neural Comput.* 12 (2000) 2305–2330.
- [78] J. Reutimann, M. Giugliano, S. Fusi, Event-driven simulation of spiking neurons with stochastic dynamics, *Neural Comput.* 15 (2003) 811–830.
- [79] G.G. Turrigiano, S.B. Nelson, Hebb and homeostasis in neuronal plasticity, *Curr. Opin. Neurobiol.* 10 (3) (2000) 358–364.
- [80] M. Tsodyks, H. Markram, The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability, *Proc. Natl. Acad. Sci. USA* 94 (1997) 719–723.
- [81] L.F. Abbott, J.A. Varela, K. Sen, S.B. Nelson, Synaptic depression and cortical gain control, *Science* 275 (1997) 220–224.
- [82] M. Griniasty, M. Tsodyks, D.J. Amit, Conversion of temporal correlations between stimuli to spatial correlations between attractors, *Neural Comput.* 5 (1993) 1–17.
- [83] N. Brunel, Dynamics of an attractor neural network converting temporal into spatial correlations, *Network* 5 (1994) 449–470.