

Eluding oblivion with smart stochastic selection of synaptic updates

Stefano Fusi^{a)}

*Institute of Neuroinformatics, UNIZH/ETH, Wintherturerstrasse 190, CH-8057 Zurich
and Institute for Theoretical Neuroscience, Columbia University, New York, New York 10032*

Walter Senn^{b)}

Institute of Physiology, University of Bern, Bülhplatz 5, CH-3014 Bern, Switzerland

(Received 8 May 2006; accepted 23 May 2006; published online 30 June 2006)

The variables involved in the equations that describe realistic synaptic dynamics always vary in a limited range. Their boundedness makes the synapses forgetful, not for the mere passage of time, but because new experiences overwrite old memories. The forgetting rate depends on how many synapses are modified by each new experience: many changes means fast learning and fast forgetting, whereas few changes means slow learning and long memory retention. Reducing the average number of modified synapses can extend the memory span at the price of a reduced amount of information stored when a new experience is memorized. Every trick which allows to slow down the learning process in a smart way can improve the memory performance. We review some of the tricks that allow to elude fast forgetting (oblivion). They are based on the stochastic selection of the synapses whose modifications are actually consolidated following each new experience. In practice only a randomly selected, small fraction of the synapses eligible for an update are actually modified. This allows to acquire the amount of information necessary to retrieve the memory without compromising the retention of old experiences. The fraction of modified synapses can be further reduced in a smart way by changing synapses only when it is really necessary, i.e. when the post-synaptic neuron does not respond as desired. Finally we show that such a stochastic selection emerges naturally from spike driven synaptic dynamics which read noisy pre and post-synaptic neural activities. These activities can actually be generated by a chaotic system. © 2006 American Institute of Physics. [DOI: [10.1063/1.2213587](https://doi.org/10.1063/1.2213587)]

Most cognitive functions rely on our remarkable ability to record and store experiences in our memory. Nerve cells (neurons) and the way they are linked through synaptic connections are widely believed to constitute the elementary storage devices of the brain. In the past 50 yr several models of artificial networks of neurons with plastic synapses have been proposed. They were constructed with simplified models of biological neurons and synapses and they were shown to have remarkably good memories, to the point that most of the investigators thought that the process of storing experiences in our memory was understood, at least at a theoretical level. In the last decade, however, the neuroscience theorists discovered that the memory performance of these networks is drastically reduced when more realistic synapses are considered. In particular, as soon as the strengths of the synaptic connections vary in a limited range, the number of memories that can be stored and remembered grows only as the logarithm of the number of synapses. The oldest experiences are forgotten at a rate that depends on the average number of synapses that are modified. A possible solution to this problem is to reduce as much as possible the number of synapses that are modified, for example, by consolidating the modifications of a randomly selected subset of synapses. We review the techniques that allow us to

increase the memory of artificial neural networks by slowing down the process of synaptic modification. The stochasticity enables an unbiased reduction in the number of synaptic modifications and plays a key role in memory formation.

I. INTRODUCTION

Highly simplified neural networks based on simple McCulloch-Pitts neurons¹ like perceptrons²⁻⁴ and Hopfield networks^{5,6} can store and retrieve an enormous number of memories. For a long time they have given the illusion that the theoretical problem of learning and memory could be solved with such simplified models of neurons and synapses. In these networks the neurons are simple threshold units, active when the total input generated by the other neurons is above a certain threshold, inactive otherwise. The synapses weigh the interaction between different neurons and they are modified through switch-like transitions from one stable state to another. The fact that the biological neurons and synapses are highly complex has been often regarded as an evolutionary accident. This intellectually unsatisfying scenario turned out to hide serious problems. Indeed, in all these models of synaptic plasticity, the memory capacity strongly depends on the number of stable synaptic states. For example, for both the perceptron and the Hopfield networks, the number of synaptic states that are necessary to learn properly all the memories grows with the number of memories that are to be

^{a)}Electronic mail: fusi@ini.unizh.ch; fusi@neurotheory.columbia.edu

^{b)}Electronic mail: senn@pyl.unibe.ch

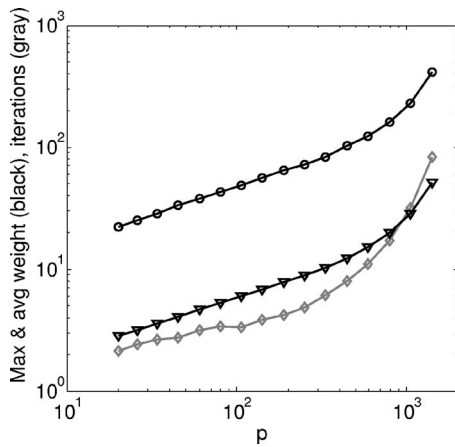


FIG. 1. Perceptron synaptic weights grow with p , the number of memories that are stored: the maximum minus the minimum synaptic weight is plotted against p (circles) in simulations of a perceptron (see, e.g., Refs. 2–4 and 7) with unbounded synapses learning random uncorrelated binary patterns ($N = 1000$ input neurons). The growth follows a power law (square root) on a wide range and then it goes up quicker as the number of memories approach the maximal storage capacity ($p = 2N = 2000$). The average absolute value of the synaptic weights grows in a similar manner (triangles). The growth is due to the increase of the number of iterations per pattern required to learn all memories (gray line and data points).

stored (see Fig. 1). This would imply that the spread between the maximum and the minimum synaptic efficacies increases with the number of memories.

Alternatively, given that biological synaptic efficacies can vary in a limited range (i.e., they are bounded), then every synapse should be modified by an amount that decreases with the number of memories that have to be stored. What happens if we do not allow for arbitrarily small steps? We will present in the next section a simple argument that shows that the memory capacity decreases so dramatically that only a few memories could be stored in the whole brain. The effects of synaptic boundaries on memories are known since the late 1980s.⁸ Such a strong reduction of memory performance was discussed already in the early 1990s^{9–11} and it is related to the entropy of the synaptic modifications following each experience to be memorized (the number of synapses that are modified and the amount by which they are modified): if large synaptic changes occur, the network is good at acquiring new memories but bad at retaining the old ones; for small synaptic modifications, old memories are overwritten at a lower rate, but less information is stored when the network goes through new experiences. The memory reduction does not depend on the specific architecture of the network, it is an inherent property of any bunch of synapses, whether they are in a network or they are considered as isolated variables.

Is it possible to reduce the learning rate as much as required to have large memory spans without reducing too much the learning step? One possibility is to reduce the learning rate by modifying only a randomly selected subset of synapses when a memory is stored.^{9–12} Such a trick reduces the effective learning rate without the recourse to unreasonably small synaptic modifications. The learning rate can be further reduced if the neural representation of memories is sparse, and if the synapses are modified only when

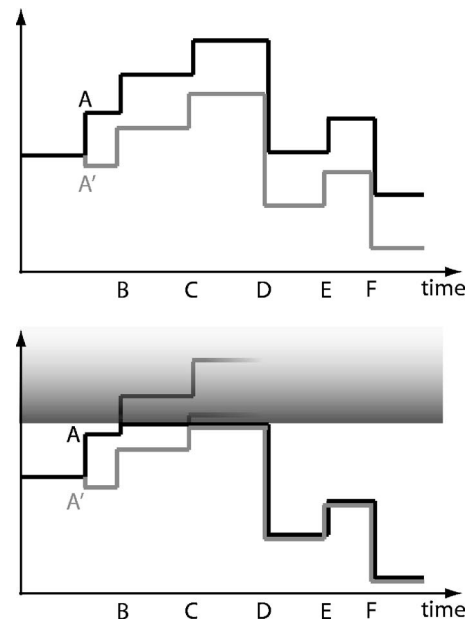


FIG. 2. Limiting the synaptic strengths between two bounds causes forgetting. Top panel: an unbounded synapse remembers about experience A. The synaptic value is plotted against time. Every experience (A–F) modifies the synapse either upward or downward. Bottom panel: the same synapse is now bounded from above (the shaded wall represents a rigid bound for the maximal value of the synapse). The trajectories that can be seen through the wall are those that the synapse would have followed, were not the bound there. Now the modifications induced by A cannot be propagated up to the present time (see the text for the explanation).

necessary, i.e., when the post-synaptic neuron does not respond as desired (see Sec. VII). We finally show that the random selection can be achieved in a natural way when the pre- and post-synaptic activity is noisy.^{13,14} This noise can be generated by a chaotic system like a network of randomly connected integrate-and-fire neurons.^{11,15}

II. LEARNING WITH REALISTIC SYNAPSES: THE PROBLEM OF OBLIVION

In a typical learning scenario, the internal state of each synapse is continuously dragged around by the stimuli corresponding to the neural representation of the experiences that are memorized. Every new experience modifies the internal state of each synapse, and it partially or totally overwrites the memory of old experiences. In the case in which each synapse is bounded, the old memory traces are not only obscured by the noise generated by the new memories, but they are actually forgotten. We illustrate this forgetting mechanism with a simple intuitive argument (see Fig. 2). We consider a learning scenario in which a synapse goes through an uninterrupted flow of experiences. Each experience is memorized by modifying the synapse. We assume that the synapse remains constant between two successive experiences. We now focus on a specific experience that occurred some time ago and we try to establish whether the synapse we are considering preserves the memory of that experience. Let us first focus on the case of unbounded synapses (Fig. 2, top panel). The synapse starts from some initial value and when it goes through experience A, it is potentiated (see the black trajectory). Successive experiences (B–F) bring the

synapse up and down, and the final, present value is determined by the sum of all these modifications. We now need a criterion to establish whether the final value is still encoding information about experience A. In order to do that, we make a virtual experiment: we go back to the past and we modify experience A in such a way that instead of potentiating the synapse, it depresses it (see the gray trajectory determined by experience A'). A necessary condition to retain memory of A is that the modification of A can be propagated through time and it can affect the final value of the synapse. If the final value is always the same whatever happened when the synapse went through experience A, then A is forgotten. In the case of unbounded synapses, the new trajectory is simply displaced by a certain amount that depends on the difference between A and A', and the final value is clearly dependent on this difference. So the final value is still correlated to the modification induced by A.

We now show that the scenario is quite different in case of bounded synapses (see Fig. 2, bottom panel). We perform exactly the same experiment as we did for the unbounded synapses. Now experience B already brings the synapse to its saturation value (indicated by a shaded wall). The synapse cannot go any farther up, so it stays at the maximum, also when experience C tries to potentiate the synapse again. When we now go back in the past and modify A into A' we see that initially the black and the gray trajectories are different, but when the synapse goes through C, it hits the upper bound and the two trajectories merge. Since then the two trajectories become identical and the final value is the same, whether the synapse went through A or A'. In this case the modification induced by A can be changed in a wide range and still it cannot affect the final synaptic value. So at the present time A has been forgotten. In this simple example we showed that forgetting is directly related to hitting the upper or the lower bound of the synapse.

How fast is such a forgetting? We can estimate the forgetting rate by computing the number of synapses that never hit one of the two boundaries. These synapses are guaranteed to retain memory of the experience we are tracking. We assume that the synaptic modifications cannot be arbitrarily small and that all the synaptic states can be visited. The probability that a synapse hits one of the two boundaries following the experience at time t is $Q(t) = p_{close}(t)p_{mod}(t)$, where $p_{close}(t)$ is the probability that the synapse is close enough to one of the two boundaries and $p_{mod}(t)$ is the probability that the synaptic modification is large enough to hit the bound. The number of synapses n that never hit a bound is

$$n = N \prod_{t=1}^p [1 - Q(t)], \tag{1}$$

where N is the total number of synapses and p is the number of experiences each synapse goes through after the memory we intend to track. If $Q_{min} = \min_t Q(t)$, then we immediately realize that n decays to zero exponentially with p :

$$n = N \prod_{t=1}^p [1 - Q(t)] < N(1 - Q_{min})^p \sim Ne^{-Q_{min}p}.$$

To retrieve information about the tracked memory, at least one synapse should be able to remember it ($n > 1$). This imposes a tight constraint on the number of patterns p that can be stored:

$$p < - \frac{\log N}{\log(1 - Q_{min})}.$$

The logarithmic dependence on N makes this limitation severe, and the neural network extremely inefficient as a memory. Most of the synaptic resources are devoted to the last stimuli seen, even when the information collected about the stimulus and stored in the synapses is more than what is needed to retrieve correctly the memorized pattern. Notice that the condition that at least one synapse remembers is a necessary condition. In general, it will not be sufficient to retrieve enough information about the memorized stimulus to reconstruct the pattern of activities imposed by the stimulus during training, not even to recognize it as familiar. p is essentially the memory span in terms of patterns whose memory trace is still in the present synaptic structure. This memory trace can be so feeble that it might be impossible to retrieve any information about the pattern of activities in the memory sliding window. This argument is very general and it can be made rigorous for practically all biologically realistic synapses.¹¹

The expression of the upper bound of p suggests a possible way out to elude the memory constraint: if consolidated synaptic changes are rare enough ($Q_{min} \ll 1$), then the sliding window in which modifications are remembered can be extended to $1/Q_{min}$ patterns:

$$p < - \frac{\log N}{\log(1 - Q_{min})} \sim \frac{\log N}{Q_{min}}.$$

The selection of the synapses that should change can be provided by the structure of the stimuli: for example, when the neural representation of the stimuli is so sparse that the synapses changed by one stimulus are unlikely to be touched by another stimulus because the patterns of activities corresponding to different stimuli have negligible overlaps. However, this is usually not enough, and additional mechanisms for further refining the selection are needed. In the absence of any information about the relevance of the stimuli, one unbiased possibility is to select randomly a fraction of the synapses which were chosen by each stimulus as candidates to change their internal state. Each synapse would change only with a given probability. This stochastic selection turns out to work nicely for random uncorrelated patterns even when the number of synaptic states is reduced to the extreme case of bistable synapses.⁹⁻¹² For example, in the case of random uncorrelated binary patterns like those used for the Hopfield model⁵ there is a minimal Q_{min} , which still allows us to acquire enough information to retrieve a pattern after a training of a single presentation. Such a Q_{min} can be reduced down to a quantity that scales like $1/\sqrt{N}$. This allows to store a number of patterns that scales like $\sqrt{N} \log N$, which is a considerable improvement. In what follows we will consider only bistable synapses since they behave qualitatively the same as multistate synapses.^{10,16} Moreover, there is accumu-

lating evidence that biological single synaptic contacts are indeed bistable on long time scales.^{17,18} Interestingly, there is also an increasing interest in implementing plastic bistable synapses in electronic hardware.^{13,19–22}

III. LEARNING ON TOP OF THE EQUILIBRIUM DISTRIBUTION

When the synapses vary in a limited range, the network tends to automatically forget for the reasons explained in the previous section. Synaptic modifications are essentially forgotten because new experiences tend to bring the distribution of the synaptic weights toward some equilibrium distribution, which depends on the statistics of the patterns and on the synaptic dynamics. This distribution does not depend on the initial condition, and hence it does not depend on the synaptic modifications that are induced by old memories. These modifications set a specific synaptic distribution, strongly correlated with the patterns to be memorized. When a specific subset of synapses is considered (for example, those synapses that are potentiated by a particular experience), the equilibrium distribution is perturbed. The successive presentations of other, dissimilar (uncorrelated) patterns tend to “levigate” the distribution and to smooth out any roughness. The distribution of synaptic weights for every subset of synapses tends exponentially to the equilibrium distribution, which means a state of oblivion. (Interestingly, oblivion comes from Latin “*ob liviscor*,” which means “towards levigating.”): whatever was the perturbation introduced by an old pattern, it is completely forgotten. To illustrate quantitatively this forgetting process, we introduce the probabilities $G_i(t)$ that synapse i is in the potentiated state at time t . Following a presentation of a pattern between time t and $t+1$, the vector G changes according to

$$G(t+1) = G(t) + Q_+(t)[\mathbf{1} - G(t)] - Q_-(t)G(t), \quad (2)$$

where $Q_+(Q_-)$ represents a vector of learning rates that encodes the probability that long term potentiation, LTP (long term depression, LTD), is induced. The product $Q_{\pm}G$ is taken component-wise. Q_{\pm} depends on the activity of the pre- and post-synaptic neurons during the presentation of the pattern and on the specific synaptic dynamics. The probability that the synapse makes a transition to the potentiated state is given by Q_+ multiplied by the probability that the synapse is actually in the depressed state $1 - G(t)$, analogously for depression. These factors reduce progressively the fraction of synapses that are allowed to change as G tends to the extremal values 0,1, and guarantees the existence of an equilibrium distribution (notice that in the case of the classic perceptron and the Hopfield model, the synapses are unbounded and there is no equilibrium distribution).

We denote the equilibrium distribution with G_{eq} and we consider the dynamics of the perturbations $G^* = G - G_{\text{eq}}$ around the equilibrium distribution. Tracing these perturbations in time and setting $\Delta G^*(t) = G^*(t+1) - G^*(t)$, we calculate:

$$\Delta G^*(t) = -Q(t)G^*(t) + L(t), \quad (3)$$

where $L(t) = (1 - G_{\text{eq}})Q_+(t) - G_{\text{eq}}Q_-(t)$ is the learning term that perturbs the equilibrium distribution, and $Q(t) = Q_+(t)$

+ $Q_-(t)$ is the rate at which the forgetting term brings G^* toward zero and, hence, G toward G_{eq} . The larger the fraction of synapses that are changed, the faster is learning, and the faster is the race toward the oblivion.

After t presentations, starting from the equilibrium distribution [$G^*(0) = 0$], we have

$$G^*(t+1) = \sum_{t'=0}^t L(t') \prod_{k=t'+1}^t [1 - Q(k)]. \quad (4)$$

The product of the $[1 - Q(k)]$ factors provokes a fast decay of the old learning terms. If the learning process is slowed down by contracting all the transition probabilities ($Q_{\pm} \rightarrow \alpha Q_{\pm}$, with $\alpha < 1$), then (4) becomes

$$G^*(t+1) = \alpha \sum_{t'=0}^t L(t') \prod_{k=t'+1}^t [1 - \alpha Q(k)],$$

where the α in front of the sum comes from the contraction of the Q_{\pm} contained in $L(t)$. If α is small:

$$G^*(t+1) > \alpha \sum_{t'=0}^t L(t') e^{-\alpha Q_{\text{max}}(t-t')}, \quad (5)$$

with $Q_{\text{max}} = \max_k Q(k)$. As α becomes small (slow learning limit), each perturbation shrinks by the same factor (the α factor in front of the sum) and the nonlinearities due to the forgetting terms disappear ($1 - \alpha Q \rightarrow 1$). The linearization is due to the fact that the forgetting term $-QG^*$ scales with the distance G^* of G from the equilibrium distribution. If this distance remains small (proportional to α), then the forgetting term scales as α^2 (both G^* and Q are proportional to α), and it can be arbitrarily small when compared to the learning term L , which instead scales linearly with α . From (5) we conclude that starting from the equilibrium distribution and appropriately reducing the learning rates Q_{\pm} (small α) will lead to a linear summation of the synaptic changes, as if synaptic saturation were not present (cf. Fig. 3). Notice that for each learning problem the learning rate α can be reduced in advance and may then remain constant throughout the learning process.

In the learning scenario considered so far the equilibrium distribution depends on the statistics of the synaptic modifications presentations, and in the slow learning limit we have

$$G_{\text{eq}} = \frac{\langle Q_+ \rangle}{\langle Q_+ \rangle + \langle Q_- \rangle},$$

where the average $\langle \dots \rangle$ is extended to all the patterns that are presented to the network (see Ref. 23 for a rigorous derivation of this formula that is valid also when the same patterns are presented more than once, as in Ref. 24).

IV. CHOOSING THE RIGHT NEURONAL THRESHOLD: THE ROLE OF INHIBITION

In the slow learning scenario (small α) the whole learning process develops in the neighborhood of the equilibrium distribution. In order to be able to read the perturbations provoked by each stimulus and to retrieve the information about its activity pattern, it is important to make the neuronal

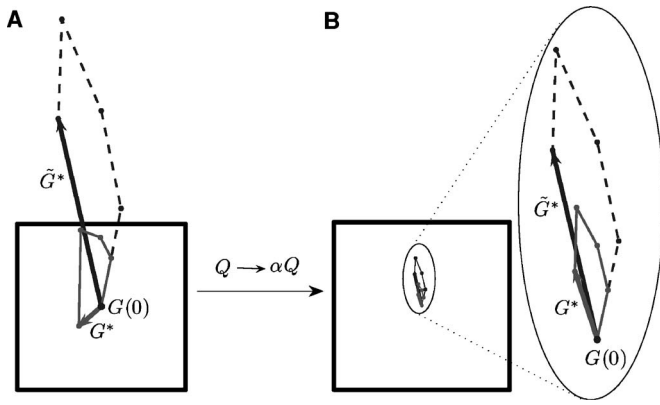


FIG. 3. Reducing the learning rates, $Q \rightarrow \alpha Q$, reduces the synaptic distortion imposed by the saturating boundary, and learning becomes possible as if the boundaries would not be present. (A) The optimal weight vector, $\tilde{G}^*(t) = \tilde{G}(t) - G(0)$, would evolve according to the dynamics $\Delta \tilde{G}^*(t) = L(t)$, which is obtained by neglecting the saturation term $-Q(t)G^*(t)$ in the dynamics for the effective weight vector $G^*(t) = G(t) - G(0)$; see Eqs. (3) and (7). Starting from $\tilde{G}^*(0) = G^*(0) = 0$ and presenting four patterns, each causing synaptic updates, the effective weight vector G^* deviates by more than 90° from the optimal weight vector \tilde{G}^* , and hence prevents a faithful learning. The arrows represent the final vectors after four iterations of the dynamics of $\tilde{G}^*(t)$ (dashed line) and $G^*(t)$ (solid line), respectively. (B) If the learning rate is reduced by a small factor α , the synaptic distortions are disproportionately reduced by a factor α^2 , while the length of each update vector is only reduced by α . This rescaling brings the effective weight vector G^* close to the optimal weight vector \tilde{G}^* . As can be proven, convergence is achieved if the effective weight vector, G^* , is defined by $G^*(t) = G(t) - g_I$, where g_I is any globally inhibitory weight vector within the synaptic bounds, for instance, with identical components $g_I = 1/2$.

dynamics sensitive to small fluctuations. For instance, if two levels of activity should be discriminated depending on the pattern to be recalled, the threshold separating the two post-synaptic current distributions should be finely tuned. Otherwise, small relative variations of the threshold (percentages of the order of α/G_{eq}) might completely disrupt the possibility to retrieve information from the synaptic matrix.

To illustrate a possible solution of this readout problem we consider a single neuron that is driven by a set of inputs ξ^μ with N components (ξ_i^μ), weighted by binary synapses of strength $J_i = 1/0$. The total excitatory current to the post-synaptic neuron is $\sum_i J_i \xi_i^\mu$. Due to the stochastic modifications triggered by presenting many different patterns, the synaptic strengths become random variables. To make the neuron sensitive to the fluctuations of the post-synaptic current around the equilibrium distribution, we subtract a negative quantity that depends on the pattern μ . For example, one might have unstructured inhibition projecting onto the same post-synaptic neuron through the inhibitory synaptic strength g_I . The total synaptic current would then be

$$h^\mu = \sum_i J_i \xi_i^\mu - g_I \sum_i \xi_i^\mu.$$

If we assume that the expected excitatory synaptic strength in the equilibrium, G_{eq} , is identical for each component (and therefore drop the index i) we can choose the global inhibitory strength as $g_I = G_{eq}$. The expected post-synaptic current in response to pattern ξ_i^μ then becomes

$$\langle h^\mu \rangle = \sum_i (G_i - G_{eq}) \xi_i^\mu = \sum_i G_i^* \xi_i^\mu.$$

A small enough threshold would be sufficient to discriminate between two distributions of h .

This subtraction has a twofold value: on the one hand it allows the neuron to read small perturbations and allows small α 's which linearize the learning process. On the other hand it permits to change only a small fraction of synapses to discriminate between patterns which should produce different responses.

A. Adapting inhibition

The problem of fine tuning is now moved to the choice of the proper level of inhibition. If the synaptic dynamics depends solely on the statistics of the patterns and it does not get any feedback from the neural dynamics (see below), then the solution of the problem requires the introduction of some additional mechanism. Interestingly this mechanism can be based on the same principles as the learning mechanism, because it is essentially required to measure the equilibrium distribution. One possibility is that the connections between the excitatory neurons of the input layer and the inhibitory cells are learned in the same way as the excitatory synaptic weights, with the only difference that the inhibitory synapses are slower (i.e. their learning rates are smaller). If the inhibitory cells are activated with the same statistics as the output excitatory cells, then the equilibrium distribution will be the same and g_I can be written as:

$$g_I = G_{eq} + g_I^*(\xi^1, \dots, \xi^p).$$

g_I^* depends on the structure of the patterns and on the details of the synaptic dynamics. If the inhibitory neurons are activated with the same statistics as the output excitatory neurons onto which they project, but they are not correlated with the specific activity, they have following each stimulus, then g_I^* will not be correlated with G^* , and it will simply represent some noise. If the activity of the inhibitory cells is correlated with the activity of the output neurons, presentation by presentation, then g_I^* will be correlated with G^* . In particular, if the excitatory connections to both the excitatory output neurons and the inhibitory neurons are updated with a Hebbian rule, then g_I^* will tend to cancel the signal produced by G^* . To preserve the memory, the signal of g_I^* should be smaller than the one produced by G^* . The only way of achieving a smaller signal without disrupting the cancellation of the equilibrium distributions is to reduce the transition probabilities of the connections between excitatory to inhibitory cells. The equilibrium distribution remains unchanged when both Q_+ and Q_- are multiplied by the same factor, but the signal depends linearly on the multiplying factor. An alternative possibility is to have an anti-Hebbian rule that updates the excitatory synapses projecting onto the inhibitory neurons.

All these mechanisms move the problem to the proper tuning of other parameters (like the transition probabilities) but prove that, in principle, it is possible to subtract the equilibrium distribution, even when the statistical structure of the patterns is not known *a priori*.

V. SENSITIVITY TO THE CODING LEVEL OF THE PATTERNS

If the patterns that are presented have very different statistics (e.g., different coding levels, or different frequency of occurrence), then different sets of stimuli might compete to bring the synaptic distribution toward completely different equilibrium distributions (see Refs. 11 and 25 for a few examples). The final equilibrium distribution will be some weighted average of these distributions and it will still allow classification. However, the rate of convergence to the final distribution can be very different from pattern to pattern, and the forgetting term will be dominated by the patterns that change the highest number of synapses. To redistribute equally the synaptic resources among the different patterns, the transition probabilities should be tuned to guarantee the same balance between the number of potentiated and depressed synapses for all patterns. For example, for uncorrelated random patterns with coding level f_μ (varying from pattern to pattern), the probability of finding a potentiating pair of pre- and post-synaptic activities is f_μ^2 and the probability of a mismatched pair of activities is $f_\mu(1-f_\mu)$. The inherent LTP transition probability q_+ should scale as $1/f_\mu^2$ and the LTD probability as $1/[f_\mu(1-f_\mu)]$ to ensure the balance for every pattern. Interestingly, in this case the equilibrium probability and the forgetting terms are the same for every pattern. Notice that balancing LTP and LTD is not necessary in case some feedback about the performance of the network (see Sec. VII).

VI. SENSITIVITY TO SIMILARITIES AND CORRELATIONS BETWEEN PATTERNS

The equilibrium distribution is also strongly biased by the correlations between patterns and by the relative frequency of presentations of different stimuli. One way to illustrate this point is to focus just on the dependency of G_{eq} on the relative frequency of presentation of each pattern. The existence of patterns that are presented more often than others would imitate the fact that highly correlated patterns are very similar and, for the purpose of illustrating our point, can be considered as repeated replicas of the same pattern. In such a case the equilibrium distribution and the ability to recall patterns change according to the relative frequency of presentation of each pattern and the transition probabilities would be weighted by this frequency. This implies that memory would be dominated by the patterns that are presented more often, while other patterns might be totally ignored, especially if they should produce responses that are in contradiction with what the majority wants. This is probably the case also in the living brain when classification is entirely free (i.e., not supervised):²⁶ subjects tend to classify patterns also according to their frequency of occurrence, especially when no feedback is provided.

VII. REFINING THE SELECTION THROUGH FEEDBACKS

Consolidated synaptic modifications should be rare events in the learning process. In the unsupervised learning scenario considered so far, the selection of the synapses that

should be changed is partially operated by the neural activity and partially by the inherent stochastic mechanism acting at the level of each synapse. Any kind of supervision might provide some feedback about the correctness of the synaptic modification and might help to refine the selection process. This further selection permits us to enlarge the memory and to partially remove the bias introduced by correlated patterns. This is how it is possible to learn to discriminate patterns that statistically would look too similar to be assigned to different classes.

A. Implementation of the feedback signal

In a supervised learning scenario the external feedback typically indicates what the desired response of the neuron is in the presence of a specific stimulus. Such feedback can be implemented by additional input currents supplied to the neurons or by any other local or global signal (e.g., neuromodulators) that modifies the way synapses are updated. These signals should affect the synaptic dynamics without changing the information to be stored. For example, they can modify the second order statistics of the trains of spikes of the pre- and post-synaptic neurons when the mean rates are the quantities to be stored.¹⁵ Alternatively, it can modify the mean firing rates in a limited range, when the quantities to be stored are just two levels of activation (for example, a high firing rate or a low firing rate^{14,27}).

In what follows we will consider a scenario in which the synapses are updated only when the response of the post-synaptic neuron does not match the one desired by the supervisor. This is the basic principle of the perceptron and the delta learning rules (see, e.g., Refs. 2–4 and 7). As soon as the response of the post-synaptic neuron is correct, the synapses are not updated. The advantage of reducing the number of synapses to be changed is twofold: (1) it slows down the forgetting process and it increases the memory span; (2) it reduces the impact of correlations between activity patterns: repeated occurrences of similar activity patterns does not change the synaptic structure any further, unless memories are not yet correctly stored.

B. The binary perceptron

The example of the perceptron with binary synapses provides a simple way to illustrate these concepts. The task is to find the synaptic weights that connect N input neurons to a single output neuron such that the total post-synaptic current h^μ is below a certain threshold θ for all patterns ξ^μ that belong to class 0 (with target output $z^\mu=0$), and above θ for all patterns of class 1 (with target output $z^\mu=1$). For the classical perceptron with analog-valued synaptic weights, an appropriate weight vector can always be found by iteratively applying the classical perceptron learning rule, provided that the patterns are linearly separable (i.e., there is a weight vector w such that $w\xi^\mu$ is a subthreshold for patterns requiring an output 0 and suprathreshold for patterns requiring an output 1). If the synaptic weights are binary, this linear separability of the two classes is not anymore sufficient to guarantee the existence of a binary separation vector. The learning problem becomes particularly hard (in fact, NP complete²⁸)

in case of online learning, in which the synapses are kept binary after every update. Recently we showed that if the number of input neurons N is large enough while the separation margin remains strictly positive (see below) then a binary weight vector that separates the classes always exists (see Refs. 29 and 30). Moreover, the perceptron learning rule will find a solution with arbitrarily high probability as N increases. However, to find such a solution, different requirements have to be met, as explained below.

1. The learning rule

The perceptron calculates the total post-synaptic current in response to a presynaptic pattern ξ^μ , $h^\mu = \sum_j (J_j - g_j) \xi_j^\mu$, and assigns 1 to the neuronal output if h^μ is above a threshold θ , and 0 if it is below θ . As introduced before, g_j represents an inhibitory synaptic strength that here is not modified. Synapses from an active presynaptic neuron can be strengthened if the target is 1, and weakened if it is 0. The synapses are only strengthened, however, if the target is 1 and the total input current is not yet too far above threshold (i.e., if $h^\mu \leq \theta + \delta$ for some margin $\delta \geq 0$). Similarly, the synapses are weakened only if the target is 0 and h^μ is not yet too far below the threshold (i.e., $h^\mu \geq \theta - \delta$). Moreover, by a *stochastic selection*, only a subset of synapses satisfying the above criteria is finally chosen to be modified. Hence, when presenting a pattern ξ^μ at time t the binary synaptic strengths $J_j (= 0, 1)$ are stochastically changed according to

$$J_j(t+1) = \begin{cases} J_j(t) + \xi_j^\mu \xi_j^\mu [1 - J_j(t)], & \text{if } z^\mu = 1, h^\mu \leq \theta + \delta, \\ J_j(t) - \xi_j^\mu \xi_j^\mu J_j(t), & \text{if } z^\mu = 0, h^\mu \geq \theta - \delta, \end{cases} \quad (6)$$

where ξ_j^\pm are binary random variables that are 1 with (small) probability Q_\pm and 0 otherwise. z^μ represents the target output for pattern ξ^μ , θ the neuronal threshold, and δ the learning margin. The factors containing J_j guarantee that a synapse will only be potentiated, provided it is currently depressed, hence the factor $[1 - J_j(t)]$, and a synapse will only be depressed provided it is currently potentiated, hence the factor $J_j(t)$. The condition on the total post-synaptic current h^μ prevents the synapse to be updated when the post-synaptic neuron responds already as desired by the supervisor. The condition on h^μ is referred to as a *stop learning*, or more appropriately as a *no-update condition*. Note that taking the expectation values in the above formula leads to the dynamics of the potentiation probabilities $G(t)$, as specified in Sec. III, Eq. (2) (with Q_\pm defined later).

2. Conditions for efficient learning

Three ingredients are necessary to guarantee the convergence to a set of synaptic weights that allows the post-synaptic neuron to respond correctly to all input patterns: (1) a global inhibition that roughly cancels the average excitation, (2) a neuronal threshold that is small compared to the total excitatory input current, and (3) a small synaptic transition probability. The basic idea is that it is possible to construct a solution in an arbitrarily small neighborhood of the initial synaptic distribution (i.e., the synaptic distribution determined by a large number of experiences that precede the

presentations of the patterns that we intend to store). Under these conditions, the departure from linear summation of the learning terms caused by synaptic saturation can be kept so small that the binary perceptron behaves like the classic perceptron with unbounded synapses. In what follows we give an intuitive argument for why the different ingredients are necessary. We will make additional assumptions to keep the argument as simple as possible. A rigorous proof of the complete perceptron convergence theorem for binary synapses can be found in Refs. 29 and 30.

3. The convergence proof

Let us assume that the components of the patterns are drawn from the same statistics, and the patterns are presented only once. In this case there is a unique equilibrium distribution G_{eq} that is the same for all components and does not change in time. We further assume that every new pattern is learned on top of the equilibrium distribution, and hence that the initial potentiation probabilities $G(0)$ are equal to G_{eq} . In the most general case, the statistics of synaptic modifications can change because the statistics of the patterns mutate, or because the external feedback changes. For the sake of simplicity, we now assume that the distribution starts from a generic $G(0)$ and that the total synaptic current is given by $h(t) = \sum_j [G(t) - G(0)] \xi_j$. Instead of adding an inhibitory term that cancels the equilibrium distribution, we subtract the initial distribution $G(0)$. We will discuss later the implications of such an assumption. Consider again $G^*(t) = G(t) - G(0)$. As in Eq. (3), we have

$$\Delta G^*(t) = -Q(t)G^*(t) + L(t), \quad (7)$$

but now $L(t) = [1 - G(0)]Q_+(t) - G(0)Q_-(t)$. Here, $Q(t) = Q_+(t) + Q_-(t)$ is the vector of effective learning rates. It is composed of $Q_+(t) = \xi^t z^t c_+(t) q$, where ξ^t is the presynaptic activity, z^t is the target output, q is a fixed transition probability determining the learning rate, and $c_+(t)$ is 1, when the post-synaptic neuron does not respond as desired by the supervisor and 0 otherwise. More precisely, we have $c_+(t) = z^t \mathcal{H}[\theta + \delta - h(t)]$, where $\mathcal{H}(\cdot)$ is the Heaviside function (1 when the argument is positive, 0 otherwise) and $h(t) = \sum_j [G_j(t) - G(0)] \xi_j^t$. Analogously, $Q_-(t) = \xi^t (1 - z^t) c_-(t) q$, where $c_-(t) = (1 - z^t) \mathcal{H}[h(t) - \theta - \delta]$.

If all components of $G(0)$ are equal to 1/2, then $L(t) = [Q_+(t) - Q_-(t)]/2$, and this ‘‘linear part’’ of the update vector is proportional to the synaptic modifications that one would have in the case of the classic perceptron with unbounded synapses. Indeed, $L(t)$ would be $+q\xi/2$ if the desired activity of the post-synaptic neuron is 1 and $h < \theta$, and $-q\xi/2$ if the desired activity is 0 and $h > \theta$. Were it not for the distortion introduced by $-Q(t)G^*(t)$, the theorem that guarantees the convergence of the perceptron learning rule for linearly separable patterns could be applied to our case. Fortunately, the distortion introduced by $-QG^*$ can be made arbitrarily small by a proper rescaling of the neuronal threshold θ and the learning rate q . Recall that forgetting term QG^* scales as α^2 , whereas $L(t)$ only scales as α , so that there is always a small α such that $L(t)$ dominates. This is also expressed by formula (5) and illustrated in Fig. 3. Notice that α is not a dynamic variable. It is just a rescaling factor that is chosen at the

beginning of the learning process, and it is never changed. The iterative process to choose its value is only virtual.

If we want to fully map the dynamics (7) onto the one of the classic perceptron, we need not only to keep the summed distortions $-QG^*$ small, but we must also ensure that the conditions for the synaptic updates, $c_{\pm}(t)$, become the same as in the classic perceptron throughout the learning process. This is also achieved by decreasing the scaling factor α . Indeed, with the scaling of the learning rate, $q \rightarrow q\alpha$, the total post-synaptic current will also be multiplied by the same factor, $h \rightarrow h\alpha$. This is because h depends on the distance between G and $G(0)$. If the threshold is similarly scaled, $\theta \rightarrow \theta\alpha$, the vanishing distortions $-QG^*$ will also lead to the same update conditions as in the classic perceptron. For small enough scaling factors α , the convergence of the learning process (7) is therefore guaranteed by the classical perceptron convergence proof.

So far we assumed that $G(0)$ was subtracted from G to compute the response of the output neuron. What happens if we decide to subtract a predefined arbitrary inhibition g_I ? Fortunately, the feedback of the supervisor would force the excitatory weight vector $G(t)$ to move toward the arbitrarily chosen inhibitory vector, even when there is no solution to the learning problem (e.g., for nonlinearly separable patterns). As soon as the components of $G(t)$ are close enough to g_I , then $G^*(t)$ can point in any direction and the solution to the learning problem can be found. As a consequence, excitation and inhibition will eventually balance. For a rigorous proof of this argument we again refer to Refs. 29 and 30.

4. Simulation results

The necessity of the three ingredients that ensure convergence is confirmed by simulations. Figure 4 shows that the number of iterations required for learning 20 random 0/1 patterns (with the same probability for 0 and 1) grows by more than a factor 100 if global inhibition (g_I) changes from a value between 0.2–0.8 to either 0.05 or 0.95. It also shows that the required iterations increase by the same factor if the threshold (θ) is beyond 0–0.2. The iterations would symmetrically increase if the threshold would be smaller than -0.2 . The span of admissible thresholds roughly corresponds to the span of admissible values for g_I , and covers roughly half the span between the maximal and minimal excitatory weight. Finally, the bottom panel shows that the transition probability Q (which is scaled by the learning rate q) needs to be small enough to ensure convergence. While no convergence would arise if Q is above some value, convergence is always possible for small enough Q , although the number of iterations decays inversely with the decreasing Q .

Figure 5 summarizes the performances for different synaptic models. The memory performance is evaluated by computing the number of patterns that can be correctly retrieved. This is done by increasing progressively the number of patterns p to be stored until the quality of retrieval (the fraction of patterns that are correctly retrieved) goes below 0.8. At the top of Fig. 5 we show the quality of retrieval versus p for the perceptron with binary synapses (crosses), the stochastic binary perceptron without the no-update condition (triangles), and the stochastic binary perceptron with the no-update con-

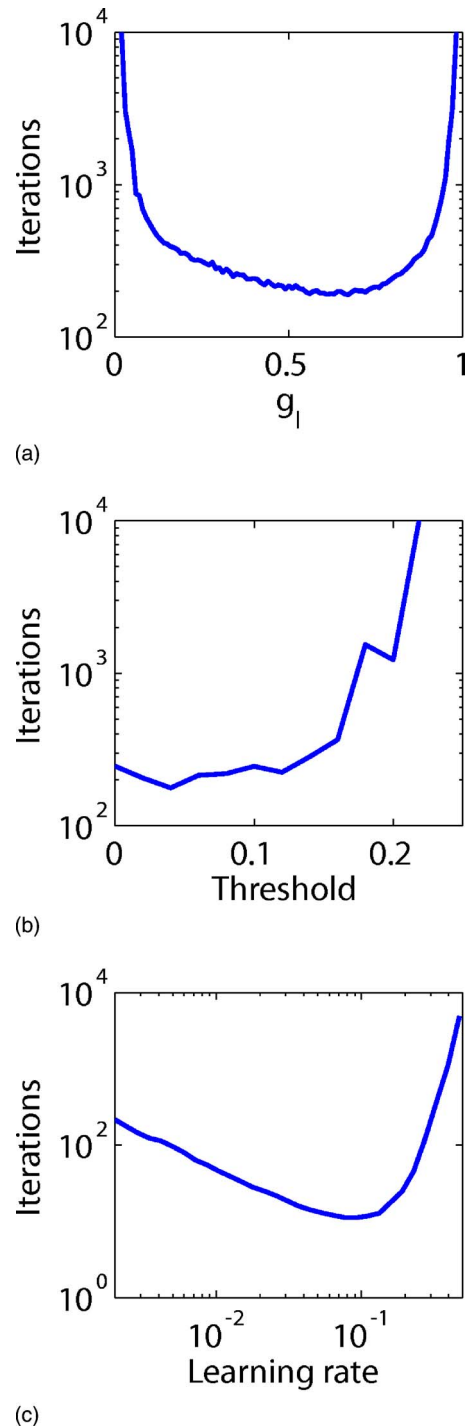


FIG. 4. Perceptron with binary synapses: required iterations to learn correctly 20 random uncorrelated patterns as a function of (a) the inhibitory synaptic strength g_I , (b) the threshold θ (scaled by $1/N$), and (c) the learning rate q . In all cases the number of input neurons was $N=500$.

dition (circles). At the bottom of Fig. 5, we show how p scales with N . Both stochastic versions of the binary perceptron perform significantly better than the deterministic binary perceptron. For the first two $p \sim \sqrt{N}$ while for the deterministic perceptron $p \sim \log(N)$. The difference in the performance becomes huge when a large number of neurons is considered.

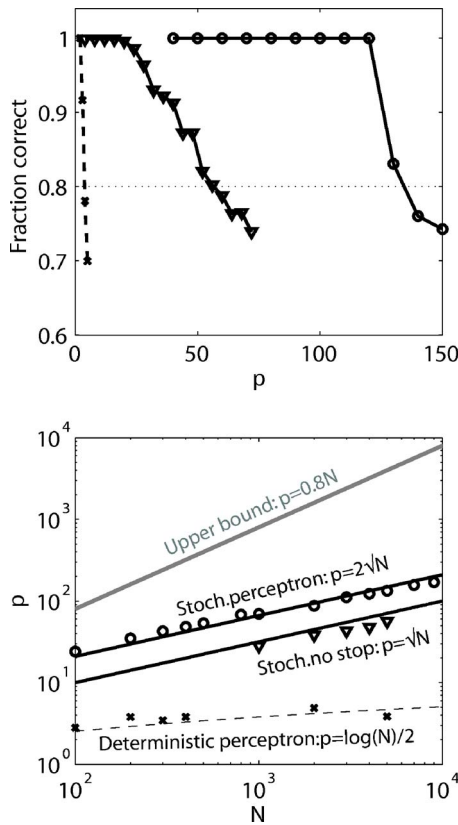


FIG. 5. Memory performance for different synaptic models. Top: the quality of retrieval (i.e., fraction of patterns that are retrieved correctly) is plotted against the total number memories that are stored during learning. As the number of memories increases, the quality of retrieval decreases because of the interference between patterns and the decay of the mnemonic trace due to the boundedness of the synapses. The quality of retrieval is plotted for three models: crosses represent the performance of a perceptron with binary synapses; triangles are for a perceptron without a no-update condition and binary synapses; and circles are for a perceptron with the no-update condition and binary synapses. In all cases the number of input neurons N is 5000 and the patterns ξ^μ to be stored (the memories) are random uncorrelated, with each component ξ_i^μ being 1 or 0 with probability 1/2. Bottom: the storage capacity—expressed as the number of patterns that can be retrieved with quality larger than 0.8—is plotted against N for the same models considered in the top panel. The gray line represents the upper bound for all networks with binary synapses (Ref. 31).

VIII. IMPLEMENTING STOCHASTIC SELECTION WITH SPIKE DRIVEN SYNAPTIC DYNAMICS

Noisy neural activity can naturally implement the stochastic selection mechanism described in the previous sec-

tions, even when the synaptic dynamics are inherently deterministic. We illustrate the main ideas by presenting a simple example. It will be clear to the reader that the same principles can be applied to a variety of synaptic models. We consider now a scenario in which the relevant information about a memory is encoded in a pattern of neural firing rates. When the plastic network goes through a new experience, each synapse memorized the activity of the pre- and post-synaptic neurons by modifying its internal state. Inspired by the cortical recordings *in vivo*, we assume that the neural activity is noisy. For example, we assume that the pre and post-synaptic neurons emit spikes according to a Poisson process with a mean frequency determined by the pattern of activities to be memorized. If these stochastic processes are statistically independent, each synapse will behave in a different way, even if it experiences the same pre- and post-synaptic average activity. The dynamics can be constructed in such a way that only a randomly chosen subset of the synapses consolidate the modifications induced by the stimulus.

We illustrate the mechanism with a specific example (see Fig. 6). We consider a synapse that is bistable in the absence of spikes, i.e., if its internal state variable X is above a threshold, then the synapse is attracted to the maximal value, otherwise it decays to the minimal value (see the midplot in Fig. 6). The maximum and the minimum are the only two stable synaptic values. Every presynaptic spike (top) kicks X up and down depending on whether the post-synaptic depolarization (bottom plots) is above or below a certain threshold. Other rules like spike timing dependent plasticity (STDP) would produce the same behavior.³²⁻³⁵ When a synapse is exposed to a stimulation, X can cross the threshold or remain on the same side as it was at the beginning of the stimulation. In the first case the synapse makes a transition to a different state (see Fig. 6 left), while in the second case, no transition occurs (see Fig. 6, right). Whether the transition occurs or not depends on the specific realization of the stochastic processes that control the pre- and the post-synaptic activity. In some cases there are enough closely spaced presynaptic spikes that coincide with elevated post-synaptic depolarization and the synapse can make a transition. In other cases the threshold is never crossed and the synapse returns to the initial value. The fraction of cases in which the synapse makes a transition determines the probability Q that

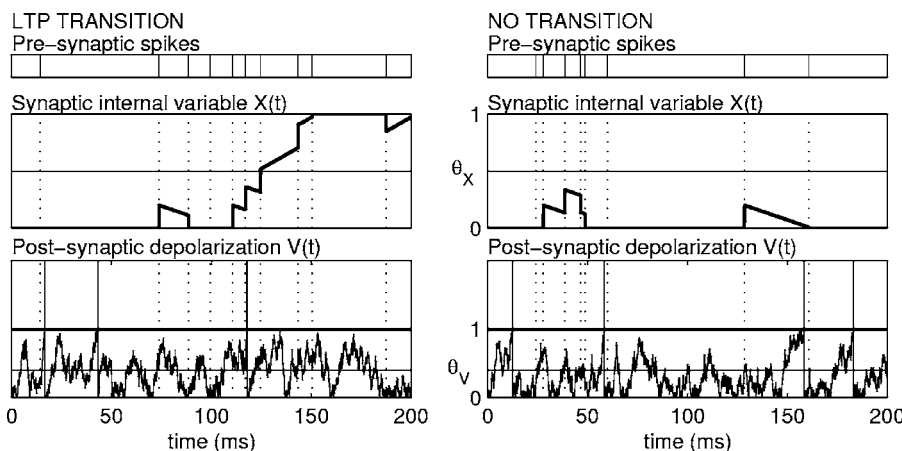


FIG. 6. Spike driven synaptic dynamics implementing stochastic selection (adapted from Ref. 10). From top to bottom: presynaptic spikes, synaptic internal state, and post-synaptic depolarization as a function of time. Left: the synapse starts from the minimum (depressed state) but between 100 and 150 ms crosses the synaptic threshold θ_x and ends up in the potentiated state. A transition occurred. Right: the synapse starts from the depressed state, it fluctuates above it, but it never crosses the synaptic threshold, ending in the depressed state. No transition occurred.

controls the learning rate. Notice that the synaptic dynamics are entirely deterministic and the load of generating stochasticity is transferred outside the synapse. Such a mechanism has been introduced in Ref. 13 and more recently has been applied to spike timing dependent plasticity,^{34,36} and to implementations of the stochastic perceptron.^{14,27} Interestingly, deterministic networks of randomly connected neurons can generate chaotic activity³⁷ and, in particular, the proper disorder that is needed to drive the stochastic selection mechanism.¹⁵

IX. DISCUSSION

Stochasticity has always played an important role in learning algorithms based on the exploration of the space of possible synaptic configurations. Reinforcement learning is certainly a well-known example.³⁸ Here we showed how stochasticity can also be important to elude oblivion in the case of realistic, bounded synapses.

Most of the classical learning algorithms assume implicitly that the variables that characterize the synaptic dynamics can vary in an unlimited range or that arbitrarily small modifications are possible. In general, as the number of memories increases, the synaptic modifications should decrease to a level that is biologically implausible. Here we showed that stochastic selection can move the problem from space—the amount by which each synapse is modified—to time—the probability that a synapse is modified, or the fraction of case in which the synapse is changed during its lifetime.

If the probability of modification can be small, then the learning rate can be reduced to very small numbers. We believe that for any physical system (including biological synapses), it is easier to reduce a probability of an event like synaptic consolidation than to regulate finely the amount by which each synapse is modified. The main reason is that small probabilities—rare events—can be readily achieved by combining independent stochastic processes. The probability of a joint event would be the multiplication of the individual probabilities and the final probability can go to zero exponentially with the number of stochastic processes that are involved in the generation of the event of synaptic consolidation. We showed in the previous section one example of a spike driven synapse implementing stochastic selection.

The drawback of the stochastic selection approach is that the learning rate is relatively low. Typically multiple presentations of the same pattern are needed in order to store enough information that the pattern can be retrieved. One of the issues that is currently being investigated is whether it is possible to speed up the learning process without sacrificing the memory span. One possibility is to combine fast and slow synapses: the first would be responsible for acquiring information, the second would retain memory of old events. A more elegant and efficient solution would be to change the probability that a synapse is consolidated according to the history of synaptic modifications (metalearning). In practice, the $Q(t)$ of Eq. (1) would decrease in time in such a way that the decay of the memory term is a power law instead of an exponential. These solutions have been recently studied in Ref. 39, and they turned out to perform orders of magnitude

better than models in which the synaptic dynamics is characterized by a single time constant (i.e., a single probability of modification).

A second issue that is still open and seems to be a general problem of all classical and more recent neural networks is the question of memory retrieval. In all the classical models and in the examples that we showed here, in order to retrieve a memory it is necessary to discriminate between very close synaptic currents: neurons that should be active receive a total synaptic current that is very close to the one that should be inactive, to a degree that again may be biologically implausible. Although it is not necessary to decide *a priori* where to place the threshold that separates active and inactive neurons, the problem of discriminating between very close values is still an open issue, also for the model proposed in the present contribution.

ACKNOWLEDGMENTS

We are grateful to Daniel D. Ben Dayan Rubin for careful reading of the manuscript, and to Daniele Fusi for the explanation of the etymology of oblivion. S.F. was supported by SNF Grant No. PP00A-106556 and by Grant No. NIH-2R01 MH58754. W.S. was supported by SNF Grant No. 3152A0-105966.

¹W. McCulloch and W. Pitts, *Bull. Math. Biophys.* **5**, 115 (1943).

²F. Rosenblatt, *Psychol. Rev.* **65**, 386 (1958); reprinted in Anderson and Rosenfeld, *Neurocomputing: Foundations of Research*, Vol. 1, pp. 89–114.

³H. Block, *Rev. Mod. Phys.* **34**, 123 (1962).

⁴M. L. Minsky and S. A. Papert, *Perceptrons* (MIT Press, Cambridge, 1969, expanded edition 1988).

⁵J. J. Hopfield, *Proc. Natl. Acad. Sci. U.S.A.* **79**, 2554 (1982).

⁶D. Amit, *Modeling Brain Function*, Cambridge University Press, Cambridge, 1989).

⁷J. Hertz, A. Krogh, and R. Palmer, *Introduction to the Theory of Neural Computation*, Addison Wesley, New York, 1991.

⁸G. Parisi, *J. Phys. A* **19**, L617 (1986).

⁹D. J. Amit and S. Fusi, *Network* **3**, 443 (1992).

¹⁰D. J. Amit and S. Fusi, *Neural Comput.* **6**, 957 (1994).

¹¹S. Fusi, *Biol. Cybern.* **87**, 459 (2002).

¹²M. Tsodyks, *Mod. Phys. Lett. B* **B4**, 713 (1990).

¹³S. Fusi, M. Annunziato, D. Badoni, A. Salamon, and D. J. Amit, *Neural Comput.* **12**, 2227 (2000).

¹⁴S. Fusi, *Rev. Neurosci.* **14**, 73 (2003).

¹⁵E. Chicca and S. Fusi, "Stochastic synaptic plasticity in deterministic aVLSI networks of spiking neurons," in *Proceedings of the World Congress on Neuroinformatics*, edited by F. Rattay (ARGESIM/ASIM Verlag, Vienna, 2001), pp. 468–477.

¹⁶S. Fusi and L. Abbott (unpublished).

¹⁷C. C. Petersen, R. C. Malenka, R. A. Nicoll, and J. J. Hopfield, *Proc. Natl. Acad. Sci. U.S.A.* **95**, 4732 (1998).

¹⁸D. H. O'Connor, G. M. Wittenberg, and S. S.-H. Wang, *J. Neurophysiol.* **94**, 1565 (2005).

¹⁹G. Indiveri, in "Advances in neural information processing systems," 2002, Vol. 15.

²⁰G. Indiveri, E. Chicca, and R. Douglas, "A VLSI array of low-power spiking neurons and bistable synapses with spike-timing dependent plasticity," in *IEEE Transactions on Neural Networks*, **17**(1), pp. 211–221, Jan. 2006.

²¹D. Badoni, M. Giulioni, V. Dante, and P. Del Giudice, "An aVLSI Recurrent Network of Spiking Neurons with Reconfigurable and Plastic Synapses," in *Proceedings of IEEE International Symposium on Circuits and Systems*, May 2006, pp. 2106–2109.

²²S. Mitra, S. Fusi, and G. Indiveri, "A VLSI spike-driven dynamic synapse which learns," in *Proceedings of IEEE International Symposium on Circuits and Systems*, pp. 2777–2780, May 2006.

²³N. Brunel, F. Carusi, and S. Fusi, *Network* **9**, 123 (1998).

- ²⁴S. Fusi, "Prototype extraction in material attractor neural networks with stochastic dynamic learning," in *Proceedings of SPIE 95*, edited by S. R. Orlando and D. Ruck, 1995, Vol. II, pp. 1027–038.
- ²⁵F. P. Battaglia and S. Fusi, *Network* **6**, 261 (1995).
- ²⁶O. Rosenthal, S. Fusi, and S. Hochstein, *Proc. Natl. Acad. Sci. U.S.A.* **98**, 4265 (2001).
- ²⁷J. Brader, W. Senn, and S. Fusi (unpublished).
- ²⁸M. Garey and D. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness* (Freeman, New York, 1999).
- ²⁹W. Senn and S. Fusi, *Phys. Rev. E* **71**, 061907 (2005).
- ³⁰W. Senn and S. Fusi, *Neural Comput.* **17**, 2106 (2005b).
- ³¹W. Krauth and M. Mezard, *J. Phys. (France)* **50**, 3057 (1989).
- ³²W. Gerstner, R. Kempter, J. van Hemmen, and H. Wagner, *Nature (London)* **383**, 76 (1996).
- ³³L. F. Abbott and S. B. Nelson, *Nat. Neurosci.* **3**, 1178 (2000).
- ³⁴N. Buchs and W. Senn, *J. Comput. Neurosci.* **13**, 167 (2002).
- ³⁵W. Senn, *Biol. Cybern.* **87**, 344 (2002).
- ³⁶W. Senn and N. Buchs, *J. Comput. Neurosci.* **14**, 119 (2003).
- ³⁷C. van Vreeswijk and H. Sompolinsky, *Neural Comput.* **10**, 1321 (1998).
- ³⁸R. Sutton and A. Barto, *Reinforcement Learning: An Introduction* (MIT Press, Cambridge, MA, 1998).
- ³⁹S. Fusi, P. J. Drew, and L. F. Abbott, *Neuron* **45**, 599 (2005).