

Learning in realistic networks of spiking neurons and spike-driven plastic synapses

Gianluigi Mongillo,¹ Emanuele Curti,² Sandro Romani³ and Daniel J. Amit^{3,4}

¹Dipartimento di Fisiologia Umana and Dottorato di ricerca in Neurofisiologia, Università di Roma La Sapienza, Rome, Italy

²Volen Center for Complex Systems, Brandeis University, Waltham, MA, USA

³INFM, Dipartimento di Fisica, Università di Roma La Sapienza, Rome, Italy

⁴Racah Institute of Physics, Hebrew University, Jerusalem

Keywords: delayed response task, long-term plasticity, persistent activity, short-term synaptic depression, working memory

Abstract

We have used simulations to study the learning dynamics of an autonomous, biologically realistic recurrent network of spiking neurons connected via plastic synapses, subjected to a stream of stimulus–delay trials, in which one of a set of stimuli is presented followed by a delay. Long-term plasticity, produced by the neural activity experienced during training, structures the network and endows it with active (working) memory, i.e. enhanced, selective delay activity for every stimulus in the training set. Short-term plasticity produces transient synaptic depression. Each stimulus used in training excites a selective subset of neurons in the network, and stimuli can share neurons (overlapping stimuli). Long-term plasticity dynamics are driven by presynaptic spikes and coincident postsynaptic depolarization; stability is ensured by a refresh mechanism. In the absence of stimulation, the acquired synaptic structure persists for a very long time. The dependence of long-term plasticity dynamics on the characteristics of the stimulus response (average emission rates, time course and synchronization), and on the single-cell emission statistics (coefficient of variation) is studied. The study clarifies the specific roles of short-term synaptic depression, NMDA receptors, stimulus representation overlaps, selective stimulation of inhibition, and spike asynchrony during stimulation. Patterns of network spiking activity before, during and after training reproduce most of the *in vivo* physiological observations in the literature.

1. Introduction

This study is the culmination of a methodological effort to capture, in a biologically realistic model, the generation of selective delay activity by repeated presentations of sequences of stimuli (Amit, 1998). The experimental motivations are animal studies which expose selective, persistent enhanced emission rates within small neural subpopulations in delayed-response tasks (e.g. Miyashita & Chang, 1988; Miyashita, 1988; Nakamura & Kubota, 1995; Erickson & Desimone, 1999). Such activity appears after several presentations of a stimulus and is not produced by novel stimuli, despite a strong selective response. The subpopulations sustaining selective delay activity for different stimuli (as many as 100) share the same set of synapses.

Modelling is in terms of neurons and synapses, which are chosen to render direct comparison with physiological experiments. Neurons are spiking elements and one can record spike rasters and spike emission statistics. Synaptic dynamics are driven by presynaptic spikes and postsynaptic depolarization (Fusi *et al.*, 2000), and can be confronted with *in vitro* experiments (e.g. see Amit & Mongillo, 2003). The choice of the plastic synapse model is guided mainly by considerations of credibility, which are unavoidable given that experimental access to *in vivo* interplay between neural and synaptic dynamics is a very remote possibility.

Models of spiking neural networks (Amit & Brunel, 1997b; Amit, 1998; Amit & Mongillo, 2003; Curti *et al.*, 2004) indicate that the observed phenomenology is reproduced by modifications of synaptic efficacies: because each selective delay population overlaps largely with the population of neurons responsive to the corresponding stimulus (Erickson & Desimone, 1999; Mongillo *et al.*, 2003), the synaptic dynamics must strengthen synapses connecting pairs of neurons responsive to a given stimulus and weaken those connecting responding to nonresponding neurons. Delay activity for stimuli in the training set emerges automatically, due to repeated presentation.

The model presented in Amit & Mongillo (2003) was limited by the following factors. (i) The stimuli were exclusive, a cell responding to, at most, one stimulus. (ii) Synaptic structuring caused an excessive increase in firing rates during stimulation, which caused instabilities in the learning process. This defect necessitated manual interventions in the simulation. (iii) Cells were linear integrate-and-fire neurons (Fusi & Mattia, 1999; Del Giudice & Mattia, 2001). In the current paper each stimulus is specified by a randomly selected set of neurons, so that a neuron can respond to more than one stimulus; see also Curti *et al.* (2004). Excessive increase in stimulus response during structuring is prevented by short-term depression of the synaptic efficacies upon activation, which is introduced as a phenomenological model (Tsodyks & Markram, 1997) with experimentally realistic time constants (S. Romani, D.J. Amit and G. Mongillo, unpublished results). The neural elements are exponential integrate-and-fire neurons, which capture experimentally observed neural response

Correspondence: Dr Gianluigi Mongillo, as ¹above.

E-mail: mongillo@surete3.romal.infn.it

Received 25 October 2004, revised 22 December 2004, accepted 10 January 2005

characteristics (Rauch *et al.*, 2003). Slow NMDA-like currents are needed to ensure the proper functioning of the network, especially to offset excessive synaptic depression immediately following the removal of a stimulus.

The model network, when subjected to repeated presentations of the stimuli in the training set, in a random sequence, autonomously develops a synaptic matrix expressing selective delay activity for each of the stimuli. Synaptic structuring occurs as a consequence of the patterns of neural activity produced by the stimuli, until a steady state for both neural activity and synaptic structuring is reached. At

asymptotic synaptic structuring, the robust behaviour of the network reproduces most of the details observed at the physiological level in delay experiments.

2. Materials and methods

2.1. The network

Table 1 gives names, symbols and values for the parameters used in this study. The network is composed of N_E excitatory and N_I inhibitory

TABLE 1. Parameters used in the simulations

Parameter		Values		
		General	Excitatory	Inhibitory
Single-cell parameters				
θ	Spike emission threshold		20 mV	20 mV
V_r	Reset potential		15 mV	10 mV
τ	Membrane time constant		20 ms	10 ms
τ_{arp}	Absolute refractory period		2 ms	2 ms
Network parameters				
c	Probability of synaptic contact	0.2		
N_E	Number of excitatory cells		8000	
N_I	Number of inhibitory cells			2000
C_E	Average number of recurrent afferent E synapses/cell (cN_E)		1600	
C_I	Average number of recurrent afferent I synapses/cell (cN_I)			400
μ_{ext}^E	Mean external current on E neurons		22.00 mV	
μ_{ext}^I	Mean external current on I neurons			18.75 mV
σ_{ext}^E	Standard deviation of external current E neurons		1.73 mV	
σ_{ext}^I	Standard deviation of external current I neurons			1.73 mV
Synaptic parameters				
J_{IE}	Synaptic efficacy E \rightarrow I	0.08 mV		
J_{EI}	Synaptic efficacy E \rightarrow I	0.18 mV		
J_{EI}	Synaptic efficacy I \rightarrow I	0.18 mV		
J_d	Depressed level of efficacy E \rightarrow E synapses	0.03 mV		
J_p	Potentiated level of efficacy E \rightarrow E synapses	0.21 mV		
γ_0	Fraction of potentiated synapses before learning	0.20		
x_{slow}^E	Fraction of slow E currents toward E neurons	0.50		
x_{slow}^I	Fraction of slow I currents toward I neurons	0.10		
τ_{slow}	Decaying time of slow E currents	100 ms		
τ_{fast}	Decaying time of fast E and I currents	0 ms		
δ	Synaptic delay	1–10 ms		
Long-term synaptic dynamics parameters				
θ_x	Threshold for synaptic transition	0.4		
θ_{LTP}	Threshold for up-regulation of X	17.5 mV		
θ_{LTD}	Threshold for down-regulation of X	15.5 mV		
α	Drift toward zero	0.0147 ms ⁻¹		
β	Drift toward one	0.0100 ms ⁻¹		
a	Amplitude of up-jump	0.25		
b	Amplitude of down-jump	0.17		
Short-term synaptic dynamics parameters				
u	Fraction of synaptic resources activated per spike	0.45		
τ_r	Recovery time of activated synaptic resources	200 ms		
Training parameters				
f	Coding level	0.15		
p	Number of stimulus prototypes in training set	7		
x_{noise}	Noise level in stimulus presentation	0, 0.1, 0.2		
T_{stim}	Duration of stimulus presentation	500 ms		
T_{delay}^E	Interval between two successive presentations	1000 ms		
C_{stim}^E	Contrast on E neurons		1.7	
C_{stim}^I	Contrast on I neurons			1.2

E, excitatory; I, inhibitory. Except for those related to short-term depression, the parameters have been chosen in the range of experimentally realistic values and have been described previously in Amit & Mongillo (2003) and Mongillo *et al.* (2003). Small adjustments have been introduced due to the population overlaps; for these we were guided by the mean-field theory of Curti *et al.* (2004). The new parameters introduced for short-term depression have been selected in the range indicated by Tsodyks & Markram (1997), guided by a mean-field theory (Romani *et al.*, 2005).

point integrate-and-fire spiking neurons, with exponential leak; the depolarization V evolves according to

$$\dot{V}(t) = -\frac{V(t)}{\tau_m} + I(t), \quad (1)$$

where τ_m is the membrane time constant and $I(t)$ is the total afferent current. Whenever the depolarization reaches a threshold θ , the cell emits a spike and remains refractory for a time τ_{arp} . Then V is reset to V_r and normal dynamics resume.

Individual postsynaptic currents obey

$$\tau_s \dot{I}_s(t) = -I_s(t) + x(t - \delta_s) \cdot J \sum_k \delta(t - t_k - \delta_s) \quad (2)$$

where τ_s is the decay time constant of channel type s , $x(t)$ is the instantaneous fraction of available synaptic resources (see below), J is the total efficacy of the synapse, t_k is the time of synaptic activation due to the k -th presynaptic spike and δ_s is a transmission delay. The dependence on the neurotransmitter involved is handled in a simplified way: fast currents, associated with AMPA receptors (excitatory) and GABA receptors (inhibitory), are taken to be instantaneous, i.e. $\tau_s = 0$. In this case, Eqn 2 becomes

$$I_s(t) = x(t - \delta_s) \cdot J \sum_k \delta(t - t_k - \delta_s).$$

Slow currents, associated with NMDA receptors (excitatory), have $\tau_s = 100$ ms. The nonlinear voltage-dependence of NMDA kinetics is not modelled. Thus, the recurrent excitatory currents have a fraction x_{slow} of slow-decaying components and the rest are instantaneous components; the recurrent inhibitory currents are instantaneous. Finite time constants for the fast receptors were introduced in testing the robustness of the system.

Each neuron also receives an excitatory, instantaneous current $I_{\text{ext}}(t)$ from outside the network (Amit & Brunel, 1997b), modelled as a Gaussian input of preassigned mean $\mu_{\text{ext}}^E, \mu_{\text{ext}}^I$ and variance $(\sigma_{\text{ext}}^E)^2, (\sigma_{\text{ext}}^I)^2$, per unit time, respectively, for excitatory and inhibitory neurons. The total current afferent on a neuron, $I(t)$ in Eqn 1, is the sum of the different contributions each evolving with its own time constant, i.e.

$$I(t) = \sum_{j,s} I_{j,s}(t) + I_{\text{ext}}(t), \quad (3)$$

where the sum on j, s is over all presynaptic neurons and all types of channels relevant for the particular synapse.

Network connectivity is random: the direct afferent presynaptic cells of a given neuron are selected, independently and randomly, by a binary process with probability c , so that each neuron receives, on average, $cN_E = C_E$ local excitatory and $cN_I = C_I$ local inhibitory recurrent contacts. Self-connection is excluded. The structure of the connectivity remains fixed throughout the simulation. The efficacies of existing excitatory as well as inhibitory synapses onto inhibitory neurons are assigned a uniform value. The recurrent excitatory-to-excitatory synapses have two possible efficacy states, potentiated J_p and depressed J_d . Prior to training, the distribution of excitatory–excitatory synapses is generated by setting each existing synapse in the potentiated (depressed) state, randomly and independently, with probability γ_0 ($1 - \gamma_0$). The synaptic delays are uniformly distributed within a bounded interval.

The uniformity in synaptic efficacies has been introduced for computational convenience to reduce memory requirements and

simulation time. It does not imply, however, uniform afferent currents to the postsynaptic neurons because of the randomness in the connectivity. The robustness of the network dynamics to the elimination of this simplifying assumption has been tested (see subsection on Robustness in Discussion).

2.2. Short-term synaptic dynamics

Following Tsodyks & Markram (1997) and Tsodyks *et al.* (1998), the synaptic connection is characterized by a given amount of ‘resource’, partitioned into three states: effective, inactive and available. Upon presynaptic emission, a fraction u of the available resource is activated, becoming effective, and then inactivated within a few milliseconds. Synaptic resources then recover to the available state, with a time constant of the order of hundreds of milliseconds (Tsodyks & Markram, 1997). Because the inactivation time is much shorter than the recovery time, the kinetics of the fraction of resources in each of the three states simplifies to the evolution of a single variable $x(t)$, the fraction of available synaptic resources at time t (Tsodyks *et al.*, 1998). The remaining equation is

$$\dot{x}(t) = \frac{1 - x(t)}{\tau_r} - ux(t) \sum_k \delta(t - t_k), \quad (4)$$

where τ_r is the time constant for resource recovery, u is the fraction of the available resources activated upon presynaptic emission and t_k is the time at which the presynaptic neuron emits spikes. The current afferent on the postsynaptic cell, via the synapse, is given by

$$I(t) = x(t) \cdot J \sum_k \delta(t - t_k) \quad (5)$$

where J (which includes a factor of u), the total synaptic efficacy, is the variation of the postsynaptic depolarization per presynaptic spike at full availability of resources, i.e. $x(t) = 1$.

Such a mechanism produces rate-dependent short-term synaptic depression because, upon arrival of a spike, the available resources decrease (Eqn 4). In between spikes, they recover to the full value, $x = 1$, on a time scale τ_r . At a low rate ($\ll 1/\tau_r$), the arriving spike finds all resources available at the synaptic site, producing maximal current (Eqn 5). As the emission rate increases above $\approx 1/\tau_r$, $x(t)$ cannot fully recover, and the current transmitted by the next spike is reduced.

2.3. Long-term synaptic dynamics

The model of the plastic synapse is characterized by an internal analogue variable $X \in [0,1]$, and by a two-state value for its stable efficacy J_p, J_d ($< J_p$) (Del Giudice *et al.*, 1998; Fusi *et al.*, 2000; Amit & Mongillo, 2003). When $X > \theta_X$, the synaptic efficacy is J_p ; for $X < \theta_X$ it is J_d . If X crosses from below to above, the result is long-term potentiation (LTP; $J_d \rightarrow J_p$); if X crosses from above to below the result is long-term depression (LTD; $J_p \rightarrow J_d$). $X = 0,1$ are reflecting barriers for the dynamics of X . These dynamics are

$$\dot{X}(t) = R(t) + H(t), \quad (6)$$

where $R(t)$ is a refresh term chosen to be

$$R(t) = -\alpha\Theta(-X + \theta_X) + \beta\Theta(X - \theta_X), \quad (7)$$

where $\Theta(x) = 1$ for $x > 0$ and 0 otherwise. $H(t)$ relates the synaptic dynamics to the pre- and postsynaptic neural activities and is responsible for synaptic transitions. It is chosen to be

$$H(t) = \sum_k F[V_{\text{post}}(t)]\delta(t - t_k^{\text{pre}}), \quad (8)$$

i.e. $H(t)$ is different from zero only at presynaptic emission and, in this case, its value depends on the instantaneous level of depolarization of the postsynaptic neuron, $V_{\text{post}}(t_k^{\text{pre}})$ through $F[\cdot]$. In our case,

$$F[V_{\text{post}}(t)] = \begin{cases} a & \theta_{\text{LTP}} \leq V_{\text{post}}(t) \leq \theta \\ -b & V_{\text{post}}(t) \leq \theta_{\text{LTD}} \text{ and refractory} \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

with $\theta_{\text{LTD}} < \theta_{\text{LTP}} < \theta$, where θ is the spike emission threshold.

This synapse behaves in a Hebbian way (Fusi *et al.*, 2000; Amit & Mongillo, 2003): when both pre- and postsynaptic emission rates are high, synaptic efficacy tends to be potentiated ($J_d \rightarrow J_p$) while a synapse connecting a high-rate presynaptic neuron to a low-rate postsynaptic one tends to be depressed ($J_p \rightarrow J_d$). Both processes are stochastic, due to the stochasticity of presynaptic spiking and of the value of the postsynaptic depolarization. No change occurs if presynaptic neurons emit at a low rate.

Figure 1 presents sample synaptic dynamic trajectories, extracted from the full simulation, to exhibit the stochasticity of LTP and LTD.




FIG. 1. Stochasticity of long-term synaptic dynamics: examples of the coupled neural–synaptic dynamics from the full simulation. Time evolution of the presynaptic depolarization [$V_{\text{pre}}(t)$, top frames], the synaptic internal variable [$X(t)$, middle frames] and the postsynaptic depolarization [$V_{\text{post}}(t)$, bottom frames] in each panel. Upon presynaptic emission (dotted vertical lines), $X(t)$ was up- or down-regulated according to the instantaneous value of $V_{\text{post}}(t)$. In between spikes, it drifted linearly toward the corresponding reflecting barrier. (A) LTP transition: both pre- and postsynaptic cell were emitting at high rates (≈ 40 Hz). (B) As in A, but no LTP. (C) LTD transition: presynaptic cell emitting at high rate (≈ 40 Hz) and postsynaptic cell at low rate ($\ll 2$ Hz). (D) As in C but no LTD. Horizontal dashed line in middle frames, threshold for synaptic transitions; horizontal dotted lines in bottom frames, thresholds for up or down regulations. For both neurons the spike emission threshold was at 20 mV, hence at the top horizontal line of the frame. The format of the present representation of the plasticity process is the analogue of Fig. 1 of Fusi *et al.* (2000); here the data are taken from a full simulation.

Fig. 1A is an LTP transition: both pre- and postsynaptic neurons were emitting at high rate (≈ 40 Hz); X started below θ_X ($J = J_d$) and was above θ_X ($J = J_p$) at the end of the stimulation interval. In Fig. 1B, both neurons emitted at same mean rate as in Fig. 1A but LTP did not occur, due to the particular realisations of the presynaptic spike train and of the postsynaptic depolarization time course. Similarly, LTD is stochastic: in Fig. 1C is a typical LTD transition, in which the presynaptic cell was emitting at high rate (≈ 40 Hz) while the postsynaptic was emitting at low rate ($\ll 2$ Hz). In Fig. 1D, under conditions of parity, LTD did not occur.

2.4. Simulation process

2.4.1. Dynamics

The instantaneous state of the network was specified by: (i) $N_E + N_I$ (10 000) analogue values of the depolarization, $V_i(t)$, for excitatory and inhibitory neurons; (ii) the analogue values of the synaptic internal variables, $X_{ij}(t)$, of the plastic synapses among excitatory neurons; (iii) the long-term efficacy of the plastic synapse $J_{ij}^{\text{EE}}(t)$; and (iv) the values of the available resources per synapse $x_{ij}(t)$ among excitatory neurons. The total number of synaptic variables depends on the randomly generated connectivity (see Section 2.1). However, due to the large value of N_E (8000), the actual number of synaptic variables is close to cN_E^2 , where cN_E is the average number of excitatory synaptic contacts per neuron

($cN_E^2 = 12\,800\,000$; in the simulations it was $12\,795\,748$). Note that the $J_{ij}^{EE}(t)$'s are kept as additional variables for computational convenience.

The simulation consisted of numerical integration of the discretized dynamic equations for the membrane depolarizations [Eqn 1 for $V_i(t)$ + condition for spike emission and refractoriness], of the short-term synaptic variable [Eqn 4, for $x_{ij}(t)$] and of the long-term synaptic variables [Eqns 6, 7 and 9 for $X_{ij}(t)$ and $J_{ij}^{EE}(t)$]. The temporal step was $\Delta t = 0.10$ ms. The initial distribution of the depolarization in the network was set uniform at a subthreshold value. Spikes began to be emitted due to the external currents. The actual value had little effect on the network dynamics; the network reached the stationary state, in which all neurons emit at the same low average rate (spontaneous activity state), within short relaxation times (≈ 100 ms). The initial values of the variables for the existing plastic synapses were set up as follows: ($J_{ij} = J_p \cap X_{ij} = 1$) randomly, with probability γ_0 ; ($J_{ij} = J_d \cap X_{ij} = 0$), with probability $1 - \gamma_0$. The initial fractions of available resources x_{ij} were distributed randomly and uniformly between 0 and 1.

The depolarization of every neuron is sequentially updated. If $V_j(t + \Delta t) \geq \theta$, a spike is delivered to all neurons postsynaptic to neuron j , and V_j is reset to $V_j = V_r$ and kept fixed for τ_{arp} . The spike adds to the value of the depolarization of the postsynaptic neuron i , at time $t + \Delta t + \delta_{ij}$, the value $J_{ij}x_{ij}(t)$. The spike also decreases the amount of available synaptic resources x_{ij} according to Eqn 4. In between spikes, both V_j and x_{ij} tend deterministically and exponentially to their steady-state values (Eqns 1 and 4).

Whenever cell j emits a spike, all internal variables (X_{ij}) of synapses which have this cell as a presynaptic cell are updated. The depolarization of each excitatory postsynaptic cell is registered and the internal variable X_{ij} is varied as described in Section 2.3. In addition each X_{ij} is refreshed linearly, according to Eqn 7. If $X_{ij} > \theta_X$, $J_{ij} = J_p$ and if $X_{ij} < \theta_X$, $J_{ij} = J_d$. Because in between spikes both synaptic variables $x_{ij}(t)$ and $X_{ij}(t)$ evolve deterministically, their actual up-dating occurs only upon presynaptic spike emission (see Mattia & Del Giudice, 2000).

Two different simulation programs have been employed to check each other. One simulation program, 'Autonomous spike learning WM', is made publicly accessible at <http://titanus.romal.infn.it/sources.html>.

2.4.2. Stimuli and learning

At the start of the simulation a set of p stimuli to be learned is set up by a binary process as follows. For every stimulus a subset of cells, both excitatory and inhibitory, is selected independently and at random with probability f (coding level). This subset of cells represents the stimulus prototype. To each stimulus prototype corresponds on average a pool of $f(N_E + N_I) = 1500$ neurons. Any other stimulus prototype shares a fraction f of these neurons ($f^2(N_E + N_I) = 225$, on average), selected at random. Stimulus prototypes are kept fixed throughout the simulation, so that one can associate to each neuron a p -bit word, $\{\zeta_i^\mu\}$, $\mu = 1, \dots, p$, where $\zeta_i^\mu = 1$ if neuron i belongs to the stimulus prototype number μ , and 0 otherwise.

The neurons for which $\zeta_i^\mu = 1$ are referred to as 'selective' to stimulus μ and those for which $\zeta_i^\mu = 0$ as nonselective for stimulus μ . This *a priori* selectivity may differ from the observed *a posteriori* selectivity, i.e. the set of neurons whose emission rate increases significantly upon presentation of stimulus μ . This is mainly because of the random connectivity and the selective stimulation of inhibitory neurons. However, the average rates in these populations are simple to monitor and are a faithful parameterization of the collective states of the network.

The p stimuli to be learned are repeatedly presented to the network in a pseudo-random sequence. Blocks of p trials are set up so that in each block the p stimuli are presented in a random order, without repetition. In each trial, the stimulus selected is presented for a period T_{stim} , followed by a delay interval T_{delay} .

The actual stimulus to be presented is a noisy version of the prototype, generated from it in the following way. A neuron of the prototype belongs to the noisy version with probability $1 - x_{noise}(1-f)$, while a neuron not belonging to the prototype belongs to the noisy version with probability $f x_{noise}$. This ensures that the average number of activated cells upon stimulation remains $f(N_E + N_I)$ (Brunel *et al.*, 1998). $x_{noise} = 0$ corresponds to the presentation of the pure prototype. During presentation of the stimulus, the mean and the variance of the external currents to the selected cells are increased by contrast factors (G_{stim}^E and G_{stim}^I both > 1). This leads to a higher average emission rate in the corresponding subset of cells.

2.5. Observables: synaptic structuring and neural activity states

2.5.1. On-line

During the simulation the following information was collected.

Average rates. The population-averaged emission rate in the stimulated, selective and nonselective neural populations, along each trial in consecutive bins of 10 ms.

Spike rasters. Spike emission times of an unbiased, random sample of 10% of excitatory cells.

Fraction of potentiated synapses within functional synaptic populations. At the end of trial number k we determine the fraction of potentiated synapses between selective cells and between selective and nonselective cells, for each of the stimuli ($\mu = 1, \dots, p$) of the training set. Two p values are obtained: a p value [$\gamma_{ss}^\mu(k)$] for selective-selective cell pairs (ss), i.e. fraction of potentiated synapses for which pre- and postsynaptic neurons respond to the same stimulus, and a p value [$\gamma_{ns}^\mu(k)$] for selective-nonselective cell pairs (ns), i.e. fraction of potentiated synapses for which pre- and postsynaptic neurons respond to different stimuli.

Internal synaptic variables. $X_{ij}(t)$ and $x_{ij}(t)$ were collected in a randomly selected sample of synapses ($\approx 0.2\%$, i.e. 2.56×10^5) between selective neurons and between selective and nonselective neurons.

Synaptic transitions. Time and type (LTP or LTD) of transitions were recorded for a sample ($\approx 0.2\%$) of synapses.

2.5.2. Off-line

The data collected on-line was elaborated off-line to monitor the network in terms of the synaptic structuring it underwent and the corresponding neural dynamics it sustained. Blocks are used as a unit for monitoring the evolution of the synaptic structuring and of the emission rates. The pseudo-random protocol we adopted guarantees that, following every block, all stimuli had been presented the same number of times.

Block-averaged emission rates during the presentation and during the subsequent delay interval. We obtained these rates by further averaging the binned average rates, within the corresponding intervals

(T_{stim} and T_{delay}), and over the p presentations within the block. In each interval the first 150 ms (i.e. the first 15 bins) were discarded to avoid transients.

Block-averaged fraction of potentiated synapses within functional synaptic populations. We averaged the p values of γ_{ss}^{μ} and the p values of γ_{ns}^{μ} at the end of each block.

Short- and long-term synaptic dynamics. The recordings of X_{ij} , x_{ij} and the data on long-term synaptic transitions were used to observe the behaviour of the synaptic device ‘*in vivo*’, i.e. embedded within a network of spiking neurons. The use of ‘*in vivo*’ is metaphoric, in the spirit of Amit (1998), because it implies the sample recording of an autonomously running simulation mimicking real experimental situations.

3. Results

The results we report are from a simulation with a training set consisting of $p = 7$ stimuli and uniform coding level $f = 0.15$, i.e. every selective population consisted of exactly $f(N_E + N_I)$ neurons. No noise was present in the stimulus presentation. The number of stimuli p was chosen low to limit the duration of the simulations; f was set high and uniform to have stable working memory (WM) activity with the ratio of potentiated : depressed efficacy not too large and to avoid finite-size effects (Curti *et al.*, 2004), given the relatively low number of neurons (10 000). The detailed description of the phenomenology in this constrained network is followed by tests of the robustness of the synaptic structuring and WM functioning to the lifting of the constraints.

3.1. Synaptic structuring

3.1.1. Basic phenomenology

In Fig. 2 we show the evolution, with the number of trials, of the fraction, γ_{ss}^{μ} , of potentiated synapses among all neurons responsive to one given stimulus (stimulus number μ) of the set, and the fraction, γ_{ns}^{μ} , of potentiated synapses between neurons responsive to

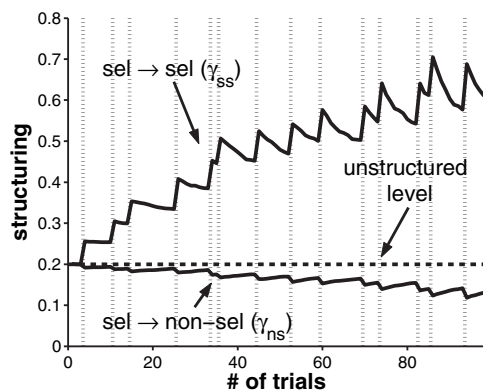


FIG. 2. Synaptic structuring, sample stimulus: fraction of potentiated synapses within the population responsive to the stimulus (sel \rightarrow sel, γ_{ss}^{μ}) and between that population and all other neurons (sel \rightarrow non-sel, γ_{ns}^{μ}) vs. the number of trials. Horizontal dashed line, initial uniform potentiation level; shaded columns, presentation of selected stimulus: γ_{ss}^{μ} increased and γ_{ns}^{μ} decreased. In between, i.e. during the presentation of other stimuli (total 7 stimuli), the trend was reversed. Parameters of Table 1.

this stimulus and neurons not responsive to it. As training proceeded, structuring took place, i.e. γ_{ss}^{μ} and γ_{ns}^{μ} began to vary with respect to their initial (unstructured) level. Whenever stimulus μ was presented, γ_{ss}^{μ} increased and, correspondingly, γ_{ns}^{μ} decreased (Fig. 2, shaded columns). However, variations in γ_{ss}^{μ} and γ_{ns}^{μ} also occurred upon presentation of any of the other stimuli of the training set: γ_{ss}^{μ} decreased while γ_{ns}^{μ} increased (Fig. 2, in between shaded columns).

The latter variations in the structuring are a consequence of the random overlaps among the neural populations coding for the stimuli, i.e. the same neuron may participate in the representation of several different stimuli, i.e. may be responsive to several stimuli. For instance, upon presentation of stimulus 1, the neurons common to stimulus 1 and stimulus 2 emitted at high rate, potentiating synapses between them. This lead to an increase in $\gamma_{ss}^{(2)}$. On the other hand, neurons selective to stimulus 2, but not to stimulus 1, emitted at low rates (during the presentation of stimulus 1), thus depressing the synapses between the neurons common to both stimuli and all other neurons in population 2. This lead to a decrease in $\gamma_{ss}^{(2)}$ which outweighed the increase mentioned above, just because there were many more neurons in population 2 which were not shared with population 1 than there were neurons belonging to both. At the same time, synapses from neurons selective to both stimuli and neurons selective only to stimulus 1 would tend to be potentiated. This would lead to an increase in $\gamma_{ns}^{(2)}$.

To summarise, the potentiation level of a given synaptic population increases (decreases) upon presentation of the corresponding stimulus, while it decreases (increases) upon presentation of the other stimuli. Consequently the structuring does not saturate, i.e. asymptotically $\gamma_{ss}^{\mu} < 1$ and $\gamma_{ns}^{\mu} > 0$. This is a fundamental difference with respect to the case of nonoverlapping stimuli. In the latter case saturation can be prevented only if both LTP and LTD probabilities in the functional synaptic population are different from zero upon presentation of a given stimulus (Amit & Mongillo, 2003).

3.1.2. Population-averaged description of synaptic structuring

As discussed in the previous section, the overlaps among selective neural populations cause variations in the synaptic structuring associated with a given stimulus upon presentation of the other stimuli (interference effect). This renders the description of the structuring process more complicated than in the case of nonoverlapping stimuli (see, e.g., Brunel *et al.*, 1998; Curti *et al.*, 2004). In the latter case, what matters is only the number of times a given stimulus was presented, because the structuring within the synaptic populations (sel \rightarrow sel and sel \rightarrow non-sel) associated with the stimulus is not affected by the presentation of other stimuli. The only interpopulation variability, at parity of number of presentations, is associated with the intrinsic stochasticity of the synaptic transitions; this variability is negligible due to the large number of synapses within each functionally homogeneous population (Amit & Mongillo, 2003).

By contrast, with overlapping stimuli the structuring of a given synaptic population depends not solely on the number of times the corresponding stimulus has been presented but on the entire presentation history. Figure 3A shows the average structuring (over the p functionally equivalent sel \rightarrow sel populations and the p equivalent sel \rightarrow non-sel populations) every p trials. The interpopulation variability was measured by the standard error of the mean (SEM) over each of the two sets of p structuring variables, and is shown in Fig. 3A as error bars. The pseudo-random presentation

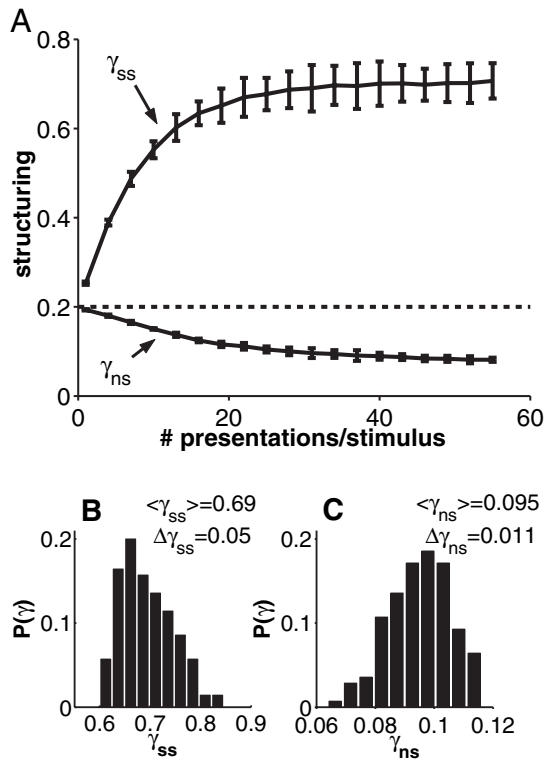


FIG. 3. Evolution of synaptic structuring. (A) The two variables characterising synaptic structuring, γ_{ss}^{μ} and γ_{ns}^{μ} , for each stimulus, averaged over all seven stimuli, vs. number of presentations/stimulus (pres/st) (curves). Error bars, interpopulation variability, SEMs over 7 values at each point. Horizontal dashed line, initial uniform potentiation level. The synaptic structuring reached asymptotically an unsaturated steady level, around which it fluctuated due to the randomness of the presentation sequence. (B) Distribution of γ_{ss}^{μ} at steady asymptotic structuring. (C) Distribution of γ_{ns}^{μ} at steady asymptotic structuring (mean and SEM given in figure). The width of these distributions is related to the LTP and LTD probabilities. See Fig. 4.

protocol (see Materials and methods) guarantees that, when averaging the structuring variables γ_{ss}^{μ} and γ_{ns}^{μ} at the end of kp trials, every stimulus would have been presented exactly k times. Thus, any residual interpopulation variability was a result of the variability in the presentation sequence.

One observes that, after ≈ 30 presentations per stimulus (pres/st), the average synaptic structuring had reached a steady, unsaturated, level ($\langle \gamma_{ss} \rangle = 0.69$ and $\langle \gamma_{ns} \rangle = 0.095$). Continuing stimulus presentation no longer affected the average synaptic structure. However, the structuring within synaptic populations corresponding to different stimuli differed, at parity of number of presentations, as shown by the error bars. The interpopulation variability was not negligible and grew with increasing number of presentations until it reached, asymptotically, a steady level. The difference in the actual structuring within a given functional population from the average was due to the fact that the number of intervening stimuli between repeated presentations of a given stimulus varied, and the interference caused by stimulus overlaps fluctuated. Figure 3B shows the distribution of γ_{ss}^{μ} at asymptotic synaptic structuring. It was obtained by collecting the fraction of potentiated synapses in all sel \rightarrow sel synaptic populations, regardless of the stimulus presented, from trial no. 217 (31 pres/st) to trial no. 385 (55 pres/st). The same was done for γ_{ns}^{μ} and the result is shown in Fig. 3C. These are steady distributions which characterise the asymptotic synaptic structuring: when stimulus μ is presented, γ_{ss}^{μ}

and γ_{ns}^{μ} are random variables drawn from the distribution in Fig. 3B and C, respectively.

The width of these distributions is related to the LTP and LTD transition probabilities per presentation; the larger the transition probabilities the larger the fluctuations (Brunel *et al.*, 1998). To show this we define q_{LTP}^{μ} as the probability that a depressed synapse undergoes LTP during the presentation of stimulus no. μ . q_{LTD}^{μ} is defined analogously. They are estimated (Amit & Mongillo, 2003) as

$$q_{LTP}^{\mu}(k) = \frac{\gamma_{ss}^{\mu}(k) - \gamma_{ss}^{\mu}(k-1)}{1 - \gamma_{ss}^{\mu}(k-1)}$$

$$q_{LTD}^{\mu}(k) = \frac{\gamma_{ns}^{\mu}(k-1) - \gamma_{ns}^{\mu}(k)}{\gamma_{ns}^{\mu}(k-1)} \quad (10)$$

where k is the trial number in which stimulus μ is presented. q_{LTP}^{μ} and q_{LTD}^{μ} are averaged in each block of p presentations, and the corresponding SEMs (over the p values in each block, separately for q_{LTP}^{μ} and q_{LTD}^{μ}) are evaluated. The result is illustrated in Fig. 4. The transition probabilities increased with the number of presentations per stimulus (from 0.07 to 0.37 for q_{LTP} and from 0.04 to 0.24 for q_{LTD}), due to the structuring process itself, until they reached a steady level. Note that, throughout the trial sequence, the probability of LTP was larger than the probability of LTD.

3.2. Neural activity

3.2.1. Response to stimulus presentation

Figure 5A shows the time course of the neural activity within the selective neural population upon presentation of the best stimulus, averaged over the first p trials, together with 10 rasters of sample selective neurons. Figure 5B shows the time course of the fraction of available synaptic resources of sel \rightarrow sel synapses over the same first block. In a given trial the $x_{ij}(t)$ of the synapses among neurons selective to the stimulus presented were sampled every 10 ms, then the average was computed over all values at the same time to give $\langle x \rangle_{ss}(t)$.

Before the presentation of the stimulus, rates were low (Fig. 5A, prestimulus interval) and synaptic resources were nearly at full availability, i.e. $\langle x \rangle_{ss} \approx 0.8$ (Fig. 5B). The abrupt increase in the

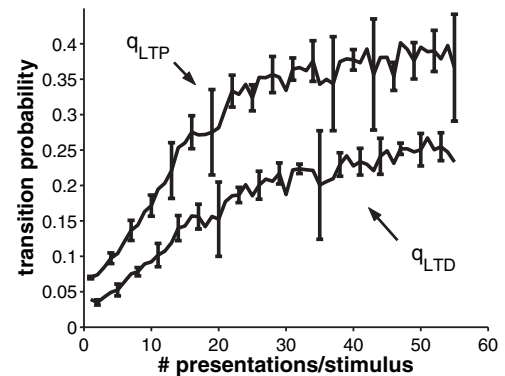


FIG. 4. Evolution of transition probabilities. Population transition probabilities averaged over all seven populations (stimuli) vs. the number of presentations. LTP probability was larger than LTD probability throughout the training stage. At asymptotic synaptic structuring, both LTP and LTD probabilities were significantly < 1 . Error bars, SEM.

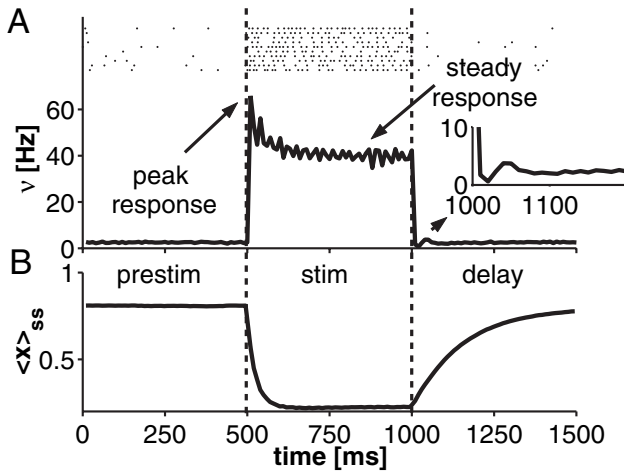


FIG. 5. Response to stimulus presentation. (A) Activity in selective population. Top, rasters of 10 sample units along the trial; bottom, activity in the selective population (in bins of 10 ms) averaged over first block. Vertical dashed lines, beginning and end of the stimulus presentation. During prestimulus interval neurons emitted at a low rate. At stimulus onset, the population activity reached a maximum (peak response) then settled at a lower level (steady response). (Inset) First 200 ms after stimulus removal. Following the removal of the stimulus, the activity dipped then overshoot and finally returned to spontaneous activity. (B) Trial-averaged time course of the fraction of available sel \rightarrow sel synaptic resources. The fast spiking during the initial transient caused a strong reduction in the fraction of available synaptic resources, with the consequent decrease in the activity level. The overshoot following the dip at the end of stimulation is due to the slow decaying NMDA-like currents (see Text). Synaptic resources recovered exponentially after stimulus removal.

external currents due to stimulus onset (see Materials and methods) caused a fraction of cells to fire almost synchronously. If that fraction was sufficiently large, its correlated firing, together with the high availability of synaptic resources, provoked further correlated firing within the population. The population activity built up in a very short time, from some 2.5 Hz to 65.8 Hz within 10 ms (peak response, Fig. 5A). The peak response, v_{peak} , was defined as the maximal population activity level (in bins of 10 ms) during stimulus presentation.

The fast rise in activity caused a temporary imbalance between inhibition and excitation. A selective neuron can fire three or four times with very short interspike intervals. Such bursting provokes the fast reduction in the fraction of available synaptic resources: $\langle x \rangle_{ss}$ decreased from 0.8 to 0.22 within ≈ 100 ms (Fig. 5B). Population activity then decreased to a steady lower level (≈ 40 Hz) as a consequence of the reduced efficacy of the recurrent excitatory synapses and of the rise in the inhibition (steady response; Fig. 5A). We defined the steady response, v_{steady} , as the average activity level in the last 350 ms of stimulus presentation.

At stimulus removal, the population activity suddenly dropped off, to a level lower than the corresponding stationary level (spontaneous or WM) due to the low fraction of available synaptic resources left over by the high rate during stimulation. However, a large amount of slow-decaying NMDA-like current had been accumulated during stimulation. This led to the overshoot (above the stationary level) in the population activity following the dip (Fig. 5A, inset). At the same time, the activity in the rest of the excitatory subnetwork, which was reduced during stimulation due to increased inhibition, rose to the spontaneous level. Note that the fraction of available synaptic resources recovered on a time scale τ_{rec} ($= 200$ ms)

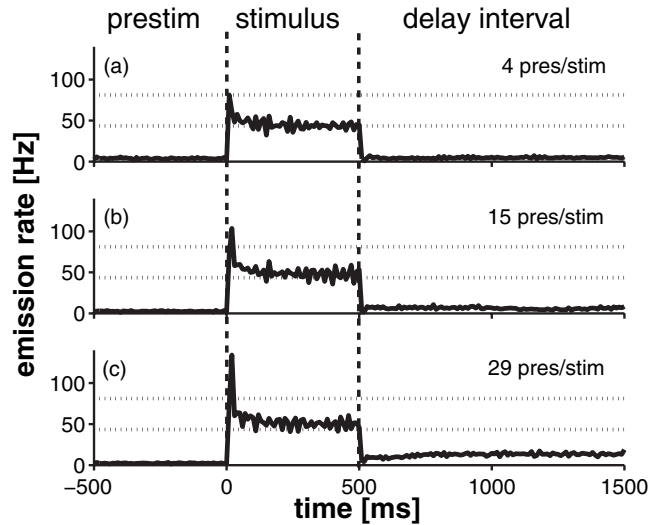


FIG. 6. Evolution of neural activity during training. The population rate (in bins of 10 ms) was sampled in three trials along the training history, after the particular stimulus has been presented 4, 15 and 29 times. Peak and steady response during stimulus presentation increased with the number of presentations. (a) 4 pres/st, no WM: following removal of stimulus, the rate in the delay interval was as in the prestimulus interval; (b) 15 pres/st, short-lived WM; (c) 29 pres/st, WM is long-lived. Vertical dashed lines, beginning and end of stimulus presentation; horizontal dotted lines, peak and steady response level in the fourth presentation.

(Fig. 5B, delay interval) longer than both the deep drop and the overshoot.

3.2.2. Basic phenomenology of neural activity

Alongside the synaptic structuring, the character of neural activity evolved as training proceeded. Samples of the evolution, for the activity within the neural population corresponding to one of the stimuli in the training set, are presented in Fig. 6. The neural emission rate was averaged over cells in the selective population, in consecutive bins of 10 ms. Following four presentations of the particular stimulus chosen, $v_{\text{peak}} = 81.2$ Hz while $v_{\text{steady}} = 43.5$ Hz (Fig. 6a). When the stimulus was novel, i.e. had been presented only a few times, the average emission rate during the delay interval (after the presentation) was as in the prestimulus interval because the synaptic strengthening was not yet sufficient to sustain reverberating activity.

As the number of presentations of a given stimulus increased, i.e. as the stimulus became familiar, the characteristics of the stimulus response and of the delay activity modified; both v_{peak} and v_{steady} increased. The increase in v_{steady} was significantly less pronounced than that in v_{peak} . For example, after 15 pres/st, the relative increase in v_{peak} was $\approx 28\%$ (to 103.6 Hz) while that in v_{steady} was 13% (to 49.0 Hz); Fig. 6b. Also, selective enhanced delay activity emerged as a consequence of the repeated presentations; Fig. 6b. At first, however, the delay activity was not very stable, i.e. it often died out before the next presentation.

Further training made the response to stimulus presentation still stronger. The relative increase in the peak response, with respect to the novel condition, was $\approx 65\%$ ($v_{\text{peak}} = 134.2$ Hz), while for v_{steady} ($= 50.4$ Hz) it was only 16%. It also rendered WM activity stable because it increased the difference between the potentiation levels in the sel \rightarrow sel and sel \rightarrow nonsel synaptic populations (see Fig. 3). After stimulus removal, the corresponding neural population reliably emitted at an enhanced rate throughout the delay interval (Fig. 6c). At

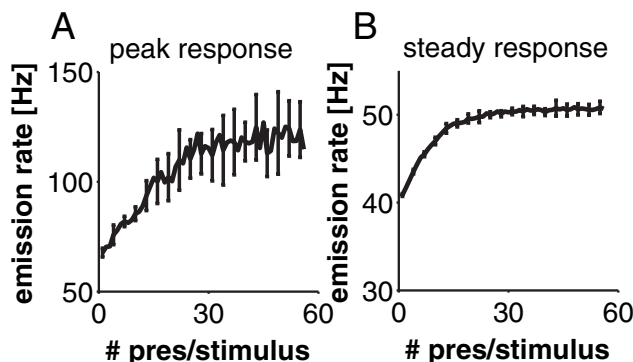


FIG. 7. Evolution of stimulus response. The peak response and the steady response were sampled in consecutive blocks, corresponding to one pres/st. The resulting p rates (one for each trial) are averaged and SEMs computed; these are the points and the error bars. (A) Peak response increased from 67.8 Hz at beginning of the training to ≈ 120 Hz at asymptotic synaptic structuring. The large increase in the error bars is related to the appearance of WM activity (see Text). (B) Steady response increased from 40.8 to 50.4 Hz within 30 pres/st. Further training did not affect the level of the steady response.

asymptotic synaptic structuring, the fraction of potentiated synapses among the neurons selective to the same stimulus was large ($\langle \gamma_{ss} \rangle \approx 0.70$), while the fraction of potentiated synapses among neurons with different selectivity was low ($\langle \gamma_{ns} \rangle \approx 0.1$). Delay emission rate in the stimulated population was ≈ 13.0 Hz, to be compared to the average emission rate within the nonstimulated populations (≈ 2.5 Hz).

Continuing stimulus presentation did not further affect the characteristics of the neural activity in the network.

3.2.3. Stimulus response evolution

The repeated presentation of a stimulus provokes synaptic potentiation within the corresponding populations and, consequently, an increase in v_{peak} and v_{steady} with increasing number of presentations. To monitor the evolution of the stimulus response, we averaged v_{peak} and v_{steady} over the p stimuli in consecutive blocks. The results are plotted against the number of presentations per stimulus in Fig. 7. The average peak response increased from 67.8 Hz at the start of training to ≈ 120 Hz (76.9%) over 30 pres/st (Fig. 7A). At the same time, the average steady response increased from 40.8 Hz to ≈ 50 Hz (22.5%; Fig. 7B). Further trials did not affect the average responses because the synaptic structuring had reached its asymptotic level (Fig. 3).

The SEMs (related to interpopulation variability) in v_{peak} and v_{steady} increased with increasing number of presentations, due to the fact that synaptic transition probabilities increased during training. Note, however, that the variability in v_{peak} was significantly larger than that in v_{steady} . This was mainly related to the appearance of WM. As the selective population sustained an enhanced rate in the delay following the presentation, the average fraction of available resources in the corresponding sel \rightarrow sel synaptic population was at a lower level than in the case in which the neural population was in spontaneous activity. For the parameters of Table 1, the WM rate was ≈ 13 Hz while spontaneous activity was at ≈ 2.5 Hz. The corresponding average fraction of available resources was $\langle x \rangle_{ss} = 0.46$ (WM) and $\langle x \rangle_{ss} = 0.82$ (spontaneous activity). Hence, when the same stimulus was presented in two successive trials, the reduced synaptic efficacy upon the second presentation made the peak response lower. With the presentation protocol chosen, stimulus repetition can occur only

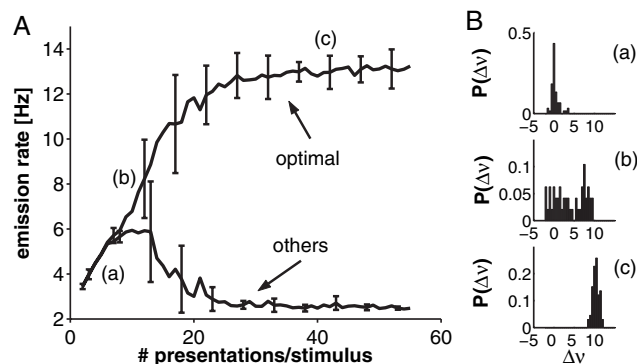


FIG. 8. Emergence of WM activity. (A) Block-averaged delay rate within the selective neural population, in optimal and other trials (see Text). Error bars, SEM. (B) Distribution of delay rate differences between optimal trials and other trials. Rate differences were collected (a) from trial no. 14 to trial no. 55; (b) from trial no. 63 to trial no. 104; (c) from trial no. 287 to trial no. 356. WM activity began to appear at ≈ 10 – 15 pres/st and stabilized at ≈ 20 pres/st. At asymptotic synaptic structuring, the average delay rate was 13.1 Hz in the optimal trials while it was 2.5 Hz in the other trials.

between two successive blocks, i.e. when the last stimulus in one block is the first stimulus in the next block. Given the low value of p , repetition occurs with relatively high probability ($1/p \approx 0.14$).

3.2.4. Emergence of WM activity

To monitor the development of selective delay (WM) activity we proceeded as follows. In each block, we collected the average delay emission rate (see Materials and methods) within the selective neural populations after the presentation of the corresponding best stimuli (optimal trials). For each selective neural population, we also collected the average delay rate in the trial successive to the optimal trial (other trials), except for the last optimal trial in the block. At the end of a block, we obtained a p vector, whose elements are the delay rates in the optimal trials, and a $(p - 1)$ vector, of the delay rates in other trials. Then we computed the average separately for the p vector, \bar{v}_s and for the $(p - 1)$ vector, \bar{v}_n , as well as the corresponding SEMs. This corresponds to measuring the average delay rate (averaged over the set of stimuli, at parity of number of presentations) in the selective neural population upon presentation of the best stimulus and upon presentation of a nonoptimal one. In this manner we tested an eventual breaking of ergodicity, namely whether the delay rate in the selective populations became distinct depending on the stimulus presented. The result is plotted against the number of presentations per stimulus in Fig. 8A.

In early stages of training, the average delay rate within a given selective neural population was independent of the stimulus presented, i.e. $\bar{v}_s \approx \bar{v}_n$ (Fig. 8A, up to 10 pres/st). There was no selective delay activity in the network. This was further checked by computing the difference between the delay rate in the selective population upon presentation of the best stimulus and the delay rate in the same population in the successive trial, i.e. upon presentation of a different stimulus. The histogram of these differences, collected from trial no. 14 to trial no. 55 (2–8 pres/st), is reported in Fig. 8, Ba. It peaks strongly around zero.

It may appear more logical to collect the delay rate before and after stimulus presentation, i.e. in the same trial, to check for the appearance of WM activity. However, because of plasticity, the presentation of a stimulus provokes strengthening of the synapses among the neurons selective to it, leading in turn to an increase in the average emission rate within the neural population. The rate difference, with respect to

the prestimulation level, increases with increasing LTP probability. In our case, from intermediate stages on, the LTP probability is relatively high (see Fig. 4). Thus, there is a potentially confounding effect in detecting the difference in emission rates related to the breaking of ergodicity (bi-stability), and not merely to the different number of presentations. Our procedure ensures that the average delay rates, upon presentation of the best stimulus and of a different stimulus, are at parity of number of presentations. Furthermore, we collected the delay rates in the selective population upon presentation of a different stimulus only in the trials successive to the optimal trials, so that potential LTD effects were negligible.

At intermediate stages (10–20 pres/st), $\tilde{\nu}_s$ and $\tilde{\nu}_n$ separated and the corresponding error bars increased (Fig. 8A). WM activity began to appear but, due to the variability in the synaptic structuring, not all populations were able to sustain stable enhanced delay activity. When the rate differences were collected from trials 63–104 (8–14 pres/st), the distribution of the differences was bimodal, i.e. in some trials there was WM activity (large delay rate differences between consecutive trials) while in others there was none (small differences). The bimodality of the distribution increased the error bars; Fig. 8, Bb). Further training rendered WM activity stable and robust and, at ≈ 20 pres/st, $\tilde{\nu}_s$ and $\tilde{\nu}_n$ were well separated and error bars shrank.

At asymptotic synaptic structuring (30–50 pres/st; see Fig. 3), both $\tilde{\nu}_s$ and $\tilde{\nu}_n$ reached a steady level where $\tilde{\nu}_s$ was significantly larger than $\tilde{\nu}_n$ (13.1 Hz vs. 2.5 Hz; Fig. 8A). All selective neural populations

exhibited enhanced delay activity following the presentation (and removal) of the corresponding stimulus, as witnessed by the distribution of rate differences collected from trial no. 287 to trial no. 356 (40–50 pres/st; Fig. 8, Bc).

3.2.5. Single-cell behaviour at asymptotic synaptic structuring

The spike rasters of a sample of 800 (10%) excitatory cells were collected upon presentation of the corresponding best stimuli over 140 consecutive trials (20 pres/st), starting at trial no. 209 (after 30 pres/st). At this stage the synaptic structuring was at its asymptotic level (see Fig. 3). Figure 9A and B shows the spike rasters together with trial averaged PSTHs for two sample cells upon presentation of the best stimulus for each.

The activity of the cells was consistent from trial to trial, i.e. the cell strongly responded to the stimulus and, after stimulus removal, exhibited enhanced delay activity, though these trials were interspersed with trials in which different stimuli were presented. A closer examination of the discharge patterns revealed significant variability from trial to trial.

As is common in experiments, to quantify this variability we estimated (from the corresponding 20 rasters) the distribution of interspike intervals (ISIs) for the two cells in Fig. 9, separately for stimulus and delay intervals. ISIs were sampled only for spikes occurring in the interval 150–500 ms from stimulus onset (stimulus) and in the interval 150–1000 ms from stimulus end (delay), where the activity of the cell is supposed to be stationary. The results are shown in

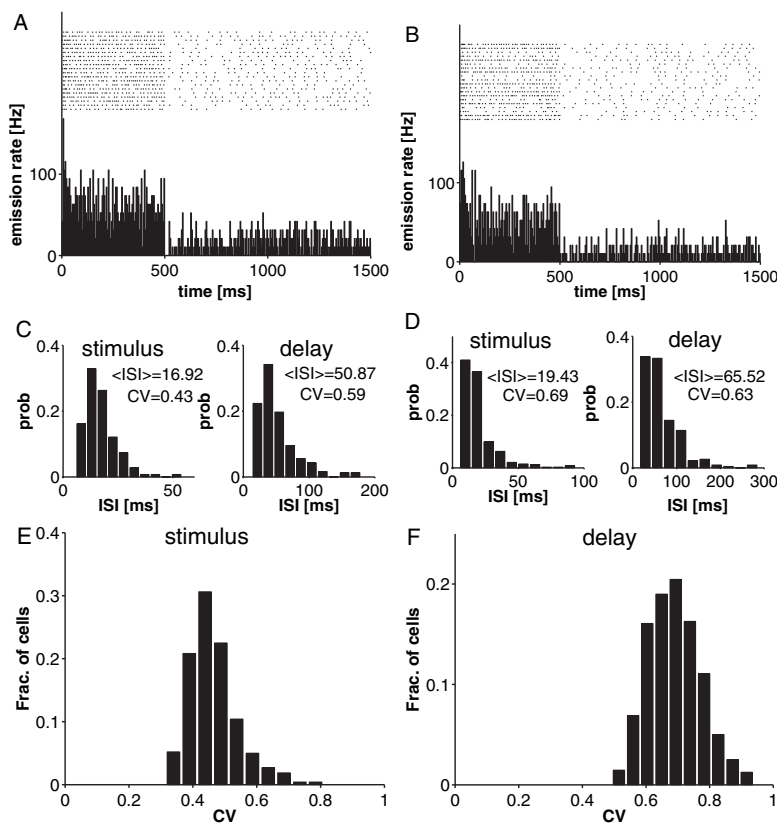


FIG. 9. Variability in single-cell discharge patterns. (A and B) Raster displays and peristimulus time histograms for two sample cells in 20 optimal trials for each cell. Both stimulus and WM activity were reproducible at rate level. Spike times varied from trial to trial. (C and D) Variability of spike times: trial-per-trial histograms of ISIs for the two sample cells, during stimulus presentation and during the subsequent delay interval. Average ISI and CV are reported in each histogram. (E and F) Distribution of CV in the neural sample (10%; 800 excitatory cells). (E) Stimulus interval, average CV 0.46; (F) delay interval, average CV 0.69.

Fig. 9C and D. The distribution of ISIs, for both cells and for both stimulus and delay intervals, is characterized by a long tail. An exponential distribution would be expected for a Poisson point process.

In Fig. 9E and F we show the distribution of the coefficient of variation (CV; the ratio between the standard deviation and the mean of the distribution of the ISIs) within the neural sample (800 excitatory cells; 10%), upon stimulus presentation (Fig. 9E) and during the subsequent delay interval (Fig. 9F). The average CV was 0.46 during stimulus presentation and 0.69 during the delay interval (WM activity). These values for CV are somewhat lower than those experimentally reported (Softky & Koch, 1993; Shadlen & Newsome, 1998), probably because of the higher emission rates in the model network; see below.

From the collected spike rasters we also extracted the single-cell average emission rate, during stimulation and in the subsequent delay interval. For each cell, the rate upon stimulus presentation was estimated by counting the total number of spikes emitted by the cell in a time window 150–500 ms from stimulus onset, in the optimal trials, and dividing by the number of trials (20) and by the time window (350 ms). The delay rate was estimated analogously in the interval 150–1000 ms following stimulus removal. The resulting distributions within the neural sample are reported in Fig. 10A for stimulus rate and in Fig. 10B for delay rate.

The rate distributions are wide, as commonly observed in experiments. The rates during stimulus presentation were in the range 11.7–81.2 Hz (average 53.0 Hz), and the delay rates 0.37–27.2 Hz (average 13.2 Hz). The variability in the average emission rate, from neuron to neuron, is mainly related to the variability in the number of connections afferent on a cell (Amit & Brunel, 1997a), the multiplicity of a cell (Curti *et al.*, 2004) and the randomness of the presentation protocol. Note, by the way, that a cell with a delay rate of 0.37 Hz has a stimulus response of ≈ 10 Hz and would appear like a cell which has no delay activity but a good stimulus response (Miyashita, 1988).

Figure 10C is the scatter plot of the emission rate during stimulus presentation vs. the emission rate during the subsequent delay interval.




FIG. 10. Distributions of single-cell emission rates. Average rates, sample as in Fig. 9. (A) During presentation of best stimulus, rates ranged from 11.7 to 81.2 Hz (average 53.0 Hz); (B) during delay interval, rates ranged from 0.37 to 27.2 Hz (average 13.2 Hz); (C) Scatter plot of rate during stimulus presentation vs. delay interval rate; each point represents a single neuron of the sample. The plot expresses strong positive correlation between rate during stimulation and rate in delay activity. Both rates are strongly affected by the (random) number of afferents.

One sees that neurons with high stimulus responses tended to have enhanced rates in the subsequent delay interval. Neurons with low stimulus responses tended to have low delay rates. This expresses the high level of correlation between the activity within the network upon stimulus presentation and in the following delay interval (retrospective activity), consistent with experiments (Erickson & Desimone, 1999).

The rate distributions measured in the simulation largely overlap with the rate distributions experimentally reported (Fuster & Alexander, 1971; Miyashita & Chang, 1988; Nakamura & Kubota, 1995; Erickson & Desimone, 1999; Naya *et al.*, 2003). However, the model network produced somewhat higher emission rates. This is likely to be related to the strong LTD, which significantly lowers the potentiation level of sel \rightarrow nonsele synaptic populations from the initial unstructured level (see Fig. 3). The average level of neural activity within the excitatory subnetwork, and consequently within the inhibitory subnetwork, decreased with training. This made stimulated and delay rates higher and nonstimulated and background rates lower. The average emission rate of nonselective cells during stimulus presentation was practically zero while during the delay interval it was ≈ 0.15 Hz. It is interesting to note that when the average emission rate of a neuron is within the experimental range the corresponding CV is also within the experimental range. This can be read from Fig. 11, which is a cell-by-cell scatter plot of the CV of the neurons of the sample vs. their average emission rate during stimulation (Fig. 11A) as well as during the subsequent delay interval (Fig. 11B). Lower, ‘biological’ emission rates correspond to CVs near 1, as would be the case for a Poisson process.

3.3. Synaptic model and life times of synaptic structure

The model of the plastic synapse is characterized by an analogue internal variable with a short time constant (≈ 20 ms), yet the internal refresh dynamics allow for the synapse to have two stable efficacy values on long time scales (Fusi *et al.*, 2000). The dynamics of the internal variable are fully controlled by the presynaptic spike events and the postsynaptic membrane depolarization. The resulting plasticity mechanism is not inconsistent with experimental findings; most of the *in vitro* stimulation protocols inducing long-term plasticity in biological synapses produce the same behaviour in the model (Amit & Mongillo, 2003). Moreover, starting from an unstructured synaptic state, the coupled neural–synaptic dynamics can drive the network into a structured state, capable of sustaining selective delay activity.

The long-term synaptic dynamics ensure that, in the absence of external stimulation, the acquired synaptic structure persists over very long time scales. This is due to the fact that, in order to have a significant probability of synaptic transitions, the average ISI must be




FIG. 11. CV vs. average emission rate. Cell-by-cell scatter plot CV vs. average emission rate. (A) Rate during stimulation by best stimulus, sample as in Fig. 9; (B) delay rate. CV is inversely correlated with the average emission rate.

of the same order as the synaptic time constant, which is the time the refresh currents take to reset the stable synaptic internal state after a jump (Fusi *et al.*, 2000; Amit & Mongillo, 2003). This occurred in a small subset of neurons (i.e. the selective neurons) upon stimulus presentation ($v_{\text{steady}} \approx 50 \text{ Hz} \rightarrow \langle \text{ISI} \rangle \approx 20 \text{ ms}$). By contrast, in the absence of external stimulation the average ISI within the excitatory subnetwork was significantly longer than the synaptic time constant ($\bar{v}_s \approx 10 \text{ Hz} \rightarrow \langle \text{ISI} \rangle \approx 100 \text{ ms} \gg 20 \text{ ms}$), leading to negligible transition probabilities.

We checked this in spontaneous and in delay activity states, in the absence of stimulation, (i) in a simulation of 200 s, with an unstructured synaptic matrix and the network in spontaneous activity ($\approx 3 \text{ Hz}$), and (ii) 20 s with the synaptic matrix at asymptotic structuring and the network in selective delay activity ($\approx 10\text{--}15 \text{ Hz}$) for one of the stimuli in the training set. We estimated the rate of synaptic transitions as the difference between the final and initial number of potentiated synapses divided by the initial number of depressed synapses and simulation duration. In both cases, we found an increase in the total number of potentiated synapses because LTP is the more probable transition, due to the lower number of jumps required (two vs. four for LTD). Transition rates were $6.1 \times 10^{-8} \text{ s}^{-1}$ for spontaneous activity and $6.4 \times 10^{-7} \text{ s}^{-1}$ during selective delay activity.

Thus, synaptic structuring is robust in the face of spontaneous and delay activity. In the absence of stimulation, synaptic transitions occur

with negligible rates. This confers stability on the learning dynamics (Del Giudice & Mattia, 2001; Amit & Mongillo, 2003).

3.4. Robustness of learning and functioning

We removed some of the constraints imposed on the simulation process. First, we allowed for random coding of the stimuli in the training set (still without noise in the stimulus presentation). The neurons belonging to a population coding for a given stimulus were randomly selected in a binary process with probability f . It resulted in variability, from stimulus to stimulus, of the total number of excitatory and inhibitory coding neurons, as well as in the relative percentage of excitatory and inhibitory neurons within a selective population (see Materials and methods). These stimuli of variable coding constituted the training set. The presentation protocol was as described in Section 2.4.2. No noticeable effects were observed, either in the average synaptic structuring or in the patterns of neural activity, throughout the course of trials (data not shown).

Next we generated the stimuli to be presented to have random coding (as above) as well as noise in their presentation. This was done by constructing each of the p stimuli in a given block by choosing one of the p prototypes (of variable coding) and suppressing the stimulation of an average fraction, x_{noise} , of the neurons of the selective population and by stimulating an equal (on average) number of excitatory cells which are not selective for that stimulus (see

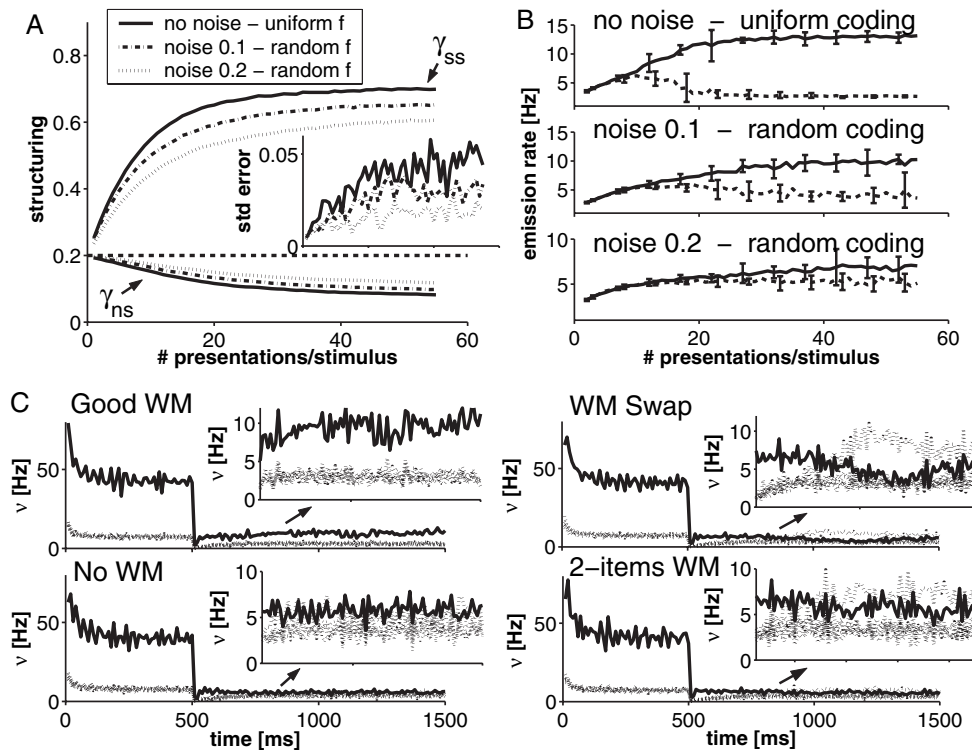


FIG. 12. Effect of noise in stimulus presentation on structuring and WM. (A) Evolution of average synaptic structuring vs. number of pres/st in the noise-free ($x_{\text{noise}} = 0$) and uniform coding level conditions (solid curves); random coding with $x_{\text{noise}} = 0.10$ (dash-dot curves), $x_{\text{noise}} = 0.20$ (dotted curves). Sel \rightarrow sel potentiation level decreased with increasing noise level; sel \rightarrow onsel potentiation level increased with increasing noise level. (Inset) Sel \rightarrow sel interpopulation variability (SEMs over the $p = 7$ values) vs. the number of pres/st (same line-pattern code). Variability in structuring decreased with increasing noise level. (B) Delay emission rate within the stimulated population in optimal trials (solid curves) and in other trials (dashed curves), (top panel) for uniform coding and $x_{\text{noise}} = 0$, clear separation of WM rate after 15 pres/st; (middle panel) random coding for $x_{\text{noise}} = 0.10$, good separation after 30 pres/st; (bottom panel) $x_{\text{noise}} = 0.20$, no satisfactory separation until 55 pres/st. (C) Neural activity in sample trials at asymptotic structuring. Solid curves, average emission rate in the selective stimulated population. Dotted curves, average emission rate in the other selective nonstimulated populations. Top left, stable selective WM activity; bottom left, no WM activity; top right, (spontaneous) transition of WM activity during the delay to a nonstimulated population; bottom right, two-item WM activity: two populations maintained elevated delay activity, the stimulated one and a random second population.

Materials and methods). The particular neurons in the subset of ‘error’ neurons varied from trial to trial. We tried $x_{\text{noise}} = 0.10$ and 0.20 .

The evolution of the synaptic structuring with the number of presentations per stimulus is reported in Fig. 12A, where we also report, for comparison, the case with uniform coding level and no noise (Fig. 3). The presence of noise had two predictable effects on the synaptic structuring. First, noise decreased the potentiation level in the sel \rightarrow sel synaptic populations and increased the potentiation level in the sel \rightarrow nonsel synaptic populations, at all stages. For $x_{\text{noise}} = 0.10$, $\gamma_{\text{ss}} = 0.65$ and $\gamma_{\text{ns}} = 0.10$, asymptotically (Fig. 12A, dash-dot curves). For $x_{\text{noise}} = 0.20$, $\gamma_{\text{ss}} = 0.60$ and $\gamma_{\text{ns}} = 0.12$ (Fig. 12A, dotted curves). These can be compared with the noise-free case where $\gamma_{\text{ss}} = 0.69$ and $\gamma_{\text{ns}} = 0.09$ (Fig. 12A, solid curves). Clearly, the probability of a sel \rightarrow sel synapse experiencing a depressing pattern of neural activity, i.e. presynaptic high rate and postsynaptic low rate, increased with increasing noise level. This kept γ_{ss} always lower than the noise-free case. Similarly, a sel \rightarrow nonsel synapse had a finite probability of experiencing a potentiating pattern of neural activity (i.e. pre- and postsynaptic rate both high), which kept γ_{ns} asymptotically larger than the noise-free case. Second, noise slows down the process of synaptic structuring (Amit & Mongillo, 2003). In the noise-free case, the potentiation level of the sel \rightarrow sel synaptic population after 15 pres/st was 88% of its asymptotic level, decreasing to 86% for $x_{\text{noise}} = 0.10$ and to 83% for $x_{\text{noise}} = 0.20$.

In the inset of Fig. 12A, we show the evolution of inter-stimulus variability among the sel \rightarrow sel structuring levels with the number of presentations per stimulus. The sel \rightarrow nonsel variability is not represented because it is significantly lower (see Fig. 3). The sel \rightarrow sel variability decreased with increasing noise level in the presented stimuli due to the reduced transition probability. Both LTP and LTD probabilities decreased because, with noise, the amplitude of the initial burst diminished (see Section 3.5). Furthermore, the LTD probability among the sel \rightarrow sel synapses was quite low because, even if some selective neurons were not stimulated, their emission rates were high due to the strong recurrent synaptic efficacy.

The functioning of the network as a WM was not disrupted by the presence of noise. This can be read from Fig. 12B, where we report the evolution, with the number of presentations per stimulus, of the average delay emission rate within the selective population in optimal trials (full curves) and in others trials (dashed curves). The appearance of stable WM activity, in the presence of noise, required a larger number of trials as a consequence of the slowing down of the synaptic structuring (see above). In the noise-free case, average delay emission rates in optimal trials and in other trials were well separated after ≈ 20 pres/st (Fig. 12B, top panel). For $x_{\text{noise}} = 0.10$, ≈ 30 – 40 pres/st were required (middle panel). Furthermore, the asymptotic delay emission rate was lower because of the lower synaptic structuring (see above). The average delay rate in optimal trials was ≈ 10 Hz (13.1 Hz for $x_{\text{noise}} = 0$) while in other trials it was ≈ 4 Hz (2.5 Hz in the noise-free case).

For $x_{\text{noise}} = 0.20$ no good separation was achieved even after 55 pres/st (bottom panel). The average delay rate in optimal trials was ≈ 7 Hz while in other trials it was ≈ 5 Hz. At this level of noise, the reliability of the network as a WM was reduced. When asymptotic synaptic structuring had been reached, the network exhibited stable, selective WM activity only in a fraction of trials. We did not attempt a quantitative estimate of the performance level. A glimpse at the observed phenomenology is presented in Fig. 12C, showing the neural activity during some sample trials (at asymptotic structuring). The top left panel shows stable selective WM activity. The selected (stimulated) population continued to emit at an

enhanced rate (solid curve) with respect to other selective populations (dotted curves), following stimulus removal, throughout the delay interval. In the bottom left panel, by contrast, after stimulus removal the average rate within the stimulated population was as in the other selective, nonstimulated populations (bottom left panel and inset). There was no WM activity. In the top right panel, during the delay interval the emission rate activity within the stimulated population went down to background level (after 500 ms from stimulus end), while the activity within one of the nonstimulated, selective populations rose up to the WM level and persisted until the end of the trial. Finally, in the bottom right panel there is an example of multiple-item WM activity (Amit *et al.*, 2003); besides the stimulated population, another nonstimulated selective population entered WM activity after stimulus removal, and both persisted throughout the delay interval.

In a separate simulation (data not shown), we added noise ($x_{\text{noise}} = 0.20$) to the stimulus presentation only at asymptotic structuring, i.e. stimulus presentation was without noise until the asymptotic synaptic structuring had been reached. In this case, selective stable WM activity was observed for 70 consecutive trials (i.e. 10 pres/st). Additional aspects of robustness are described in the Discussion.

3.5. ‘In vivo’ behaviour of plastic synapses

As observed in Section 3.1.2, the average synaptic transition probabilities per presentation grew along the course of training until they reached steady levels at asymptotic synaptic structuring. This increase was due to the fact that both v_{peak} and v_{steady} increased with structuring. As the ISIs during stimulus presentation became shorter, the transition probabilities increased.

The average ISI in a sample (10%) of selective cells upon presentation of the best stimulus was estimated in two periods during the stimulation. For each cell and in each trial, ISIs were collected over the first four spikes (three values) from stimulus onset and over the last 350 ms of stimulus presentation, uniting values for equal numbers of presentations. We considered the first four spikes because four is the minimum number of jumps required for both LTP and LTD to occur. The average ISI over the first four spikes went from ≈ 23.5 ms at the start of training to 16.0 ms, asymptotically. The average ISI during steady response went from 26.4 ms to 21.0 ms (Fig. 13A). With training, both averages decreased and, consequently, long-term synaptic transitions tended to occur with increasing probability (see Fig. 4).

During the initial transient (i.e. the first four spikes), the stimulated neurons emitted with ISIs shorter than those in the subsequent late phase throughout training history (Fig. 13A). As a consequence the synaptic transitions tended to occur earlier in the stimulation interval where the ISIs were shorter. To show this, we collect the transition times within the sel \rightarrow sel (LTP) and sel \rightarrow nonsel (LTD) synaptic populations, upon the presentation of each stimulus separately, over consecutive blocks. The transition time is the time, from stimulus onset, at which the synaptic internal variable $X(t)$ crosses the threshold (LTP from below to above, LTD from above to below), and does not return back (see Fig. 1). The stimulation interval was divided into 10 bins of 50 ms. In each bin we counted the total number of transitions that occurred in the entire block of p trials. The probability of transition in a given bin was estimated as the number of transitions in that bin divided by the total number of transitions that took place in the entire stimulation interval in the p trials. The upper and lower bounds of the confidence interval of the estimated probabilities (Meyer, 1965) are given by

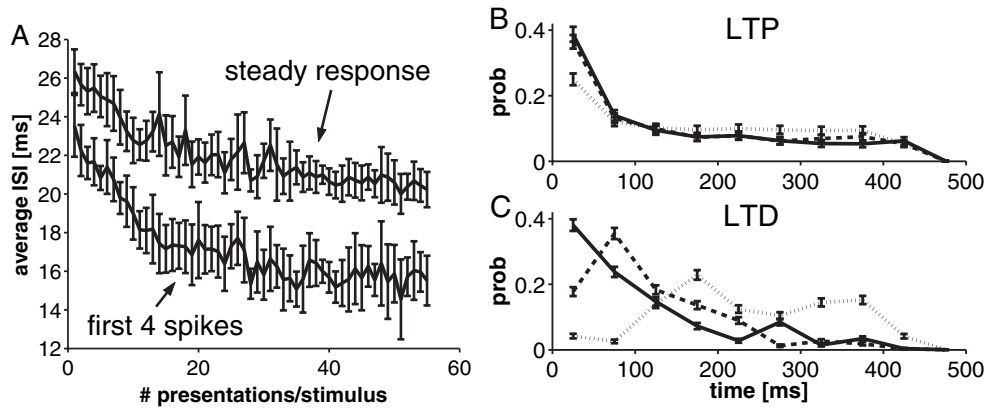


FIG. 13. Effect of training on timing of spikes and synaptic transitions during stimulation. (A) Average ISI vs. the number of pres/st: lower curve, average ISI over first four spikes from stimulus onset; upper curve, average over the last 350 ms of stimulation interval. Error bars are SEM. (B,C) Probability of synaptic transition vs. time elapsed from stimulus onset (bins, 50 ms); the distributions of times are estimated after one presentation/stimulus (dotted curves), after 10 pres/st (dashed curves), after 30 pres/st (solid curves; asymptotic). (B) LTP transitions; (C) LTD transitions. Error bars are confidence intervals (Eqn 11, $k = 2$). The neurons emitted with shorter ISIs in the early phase of stimulus response than in the late phase. Training decreased ISIs immediately after stimulus onset, increasing the probability of early transitions. LTP tended to occur early with respect LTD throughout the training, due to the smaller number of jumps required per transition (2 vs. 4).

$$P_{\text{up/low}} = \frac{Pn + k^2/2 \pm k[P(1-P)n + k^2/4]^{1/2}}{n + k^2} \quad (11)$$

where n is the total number of transitions, P is the fraction of transitions within a given bin (i.e. the estimated probability) and k ($= 2$) is the width of the confidence interval in SD. The distributions of transition occurrence times (from stimulus onset) during the first (dotted line), 10th (dashed line) and 30th (solid line) presentation are shown in Fig. 13B and C for LTP and LTD, respectively. LTP tended to occur early in the stimulation interval with respect to LTD, throughout the training (Fig. 13), because the minimal number of up-jumps required for LTP is two while for LTD at least four down-jumps are required. Later in training, the probability of ‘early’ transitions increased. The increase in the transition probability at the beginning of the stimulation interval was more evident for LTD (Fig. 13C). At the beginning of training, the probability of LTD occurring became significant after ≈ 200 ms from stimulus onset because four down-jumps are required for LTD. At asymptotic synaptic structuring, LTD occurred with elevated probability within the first 50 ms, due to the shorter ISIs (Fig. 13C, solid line).

Throughout training history, the average ISI over the first four spikes decreased more significantly than that during the late response (47 vs. 26%), mimicking the corresponding variation in v_{peak} (with respect to v_{steady}). A rough inverse proportionality between the average ISI over first four spikes and v_{peak} is indeed expected. Moreover, one also expects a correspondingly larger contribution to the variation of the transition probabilities. The larger the burst amplitude (and its variation) the larger the transition probabilities (and their variations). Interestingly, the strong dependence of the transition probabilities on the amplitude of the peak response can be also seen in Fig. 4, where a significant increase in the error bars from 15 to 20 pres/st is evident. This is related to the appearance of WM activity and its effect on the amplitude of the peak response (see Section 3.2.3).

To test the effect of the initial transient of the stimulus response on the transition probabilities, we manipulated the amplitude of the initial burst by desynchronising the external selective currents during stimulus presentation. In each trial, a delay was randomly associated with each selective neuron, representing the time from nominal stimulus onset to the instant in which the external afferent current was actually increased (see Materials and methods). The delays were

uniformly distributed between 0 and δ_{act} ms, with steps of 0.1 ms (equal to the time step used in simulation).

Figure 14A shows the time course of stimulus response (averaged over all stimuli) in an unstructured synaptic matrix for $\delta_{\text{act}} = 20$ ms (full curve), $\delta_{\text{act}} = 50$ ms (dashed curve) and, for comparison, for synchronous activation of external afferents (main simulation; thin solid curve). Desynchronising the activation of the afferents during stimulus presentation resulted in a reduction in the peak response. It also increased the time, from stimulus onset, to reach maximal emission. Both effects are related to the fact that the fraction of cells firing almost synchronously at stimulus onset decreases with increasing δ_{act} (see Section 3.2.1). The peak response at the start of training decreased by 22% from 67.8 Hz (for synchronous activation) to 55.5 Hz for $\delta_{\text{act}} = 20$ ms, and by 28% to 52.8 Hz for $\delta_{\text{act}} = 50$ ms. Correspondingly, the time to peak response increased from ≈ 10 ms (synchronous activation) to ≈ 30 ms for $\delta_{\text{act}} = 20$ ms and to ≈ 70 ms for $\delta_{\text{act}} = 50$ ms (Fig. 14A).

Figure 14B shows the evolution along the training history of the average v_{peak} (over all stimuli) for $\delta_{\text{act}} = 20$ ms (solid curve), $\delta_{\text{act}} = 50$ ms (dashed curve) and in the case of synchronous activation (thin solid curve) for comparison. For clarity we have only shown the SEMs (over the p values) for the case of asynchronous activation. The SEMs for synchronous activation can be read from Fig. 7A. With increasing δ_{act} the peak response increased more mildly along the training history (Fig. 14B). For $\delta_{\text{act}} = 20$ ms, the average peak response increased by 74% from 55.5 Hz at the start of the training to 96.9 Hz at asymptotic synaptic structuring; for $\delta_{\text{act}} = 50$ ms it increased by 33% from 52.8 Hz to 70.2 Hz. In the synchronous case the relative variation in the peak response was 76.9%.

In all cases, the structuring trajectories and, consequently, the evolution of the steady response and of the delay activity were very close despite of the significant variation in the evolution of the peak response (data not shown). Asymptotically, for $\delta_{\text{act}} = 20$ ms, $\gamma_{\text{ss}} = 0.68$ and $\gamma_{\text{ns}} = 0.09$; for $\delta_{\text{act}} = 50$ ms, $\gamma_{\text{ss}} = 0.66$ and $\gamma_{\text{ns}} = 0.09$; (in the main simulation, $\gamma_{\text{ss}} = 0.69$ and $\gamma_{\text{ns}} = 0.09$, $\delta_{\text{act}} = 0$ ms). Finally, in panels Fig. 14C and D we show the evolution, with the number of pres/st, of the average LTP and LTD transition probabilities per presentation for $\delta_{\text{act}} = 20$ ms (solid curves) and $\delta_{\text{act}} = 50$ ms (dashed curves). Again, for comparison, we show the corresponding data (thin curves) for synchronous activation (see Fig. 4

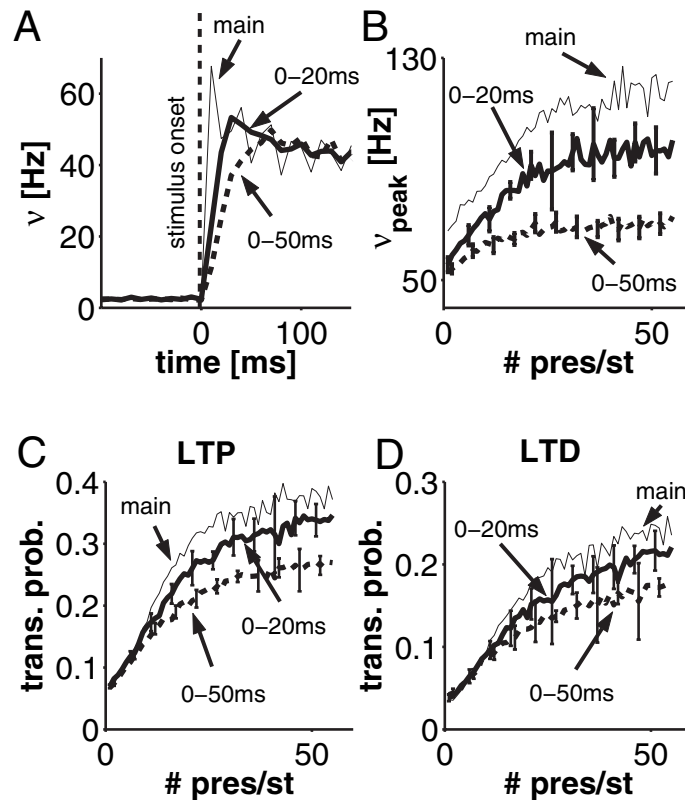


FIG. 14. Effects of desynchronization of external afferents on learning dynamics. (A) Average time course of stimulus response at unstructured synaptic matrix (before training), for asynchronous stimulation (solid curve, $\delta_{\text{act}} = 20$ ms; dashed curve, $\delta_{\text{act}} = 50$ ms) and synchronous activation (thin solid curve, $\delta_{\text{act}} = 0$ ms). Asynchronous activation decreased the burst amplitude and increased the time to peak response. (B) Average peak response vs. the number of pres/st. Asynchronous activation reduced the variation in the peak response through training history. (C) Average LTP probability vs. number of pres/st. (D) Average LTD probability vs. number of pres/st. Asynchronous activation reduced the variation of the transition probabilities along the training. SEMs are reported, as error bars, only for asynchronous activation.

for SEMs). In all cases the transition probabilities increased with the number of presentations per stimulus, and the LTP probability was larger than the LTD probability throughout the training. However, in the case of asynchronous activation the increase in the transition probabilities decreased with increasing δ_{act} : q_{LTP} went from ≈ 0.07 to 0.33 for $\delta_{\text{act}} = 20$ ms and to 0.26 for $\delta_{\text{act}} = 50$ ms (0.37 in the synchronous case). q_{LTD} increased from ≈ 0.04 to 0.21 for $\delta_{\text{act}} = 20$ ms and to 0.16 for $\delta_{\text{act}} = 50$ ms (0.24 in the synchronous case).

4. Discussion

The main result achieved in the present work is the demonstration, by simulations, that in the large space of parameters characterising the ensemble of simple, universal microscopic elements, i.e. neurons and synapses, a suitable zone can be found in which the network functions as a system which dynamically stores and recalls a set of randomly chosen stimuli. The appropriate zone of parameters is within ranges of biological credibility. Moreover, the characteristics of the neural dynamics express most of the details of the corresponding observed physiological phenomena.

The network is subjected to a stream of stimulus-delay trials in which one of the stimuli belonging to the training set is presented, followed by a delay during which no stimulus is presented. The patterns of neural activity induced by the repeated presentation of the set of stimuli, via spike-driven local synaptic plasticity dynamics, lead to synaptic structuring. During stimulus presentation, the concurrent activation at high rates of the cells coding for the stimulus increases the

potentiation level (fraction of synapses in the potentiated state) of the synaptic population connecting the neurons selective for the stimulus. At the same time, the synaptic dynamics decrease the potentiation level of the synaptic population from stimulated to nonstimulated neurons. When the difference between the potentiation levels reaches a suitably high level, the neural population becomes capable of reliably sustaining reverberating activity in the absence of external selective inputs (Brunel *et al.*, 1998; Curti *et al.*, 2004).

No external intervention or artificial stop-learning conditions are involved at any stage. After sufficiently long training (30–35 pres/st; note that WM appears much earlier in the course of training; see Section 3.2.4), the coupled neural–synaptic dynamics reach a stable, asymptotic configuration. Neural activity, whether the network is stimulated (by familiar stimuli) or not, no longer affects the synaptic structuring itself (in a statistical sense). Consequently, the patterns of neural activity exhibited by the network remain stable apart from fluctuations, related to the corresponding fluctuations in the synaptic structuring, which do not affect the qualitative functioning of the network. If the training set and the frequency with which stimuli are presented (presentation rate) remain unchanged, both the synaptic structuring and the corresponding patterns of neural activity persist.

A particularly interesting feature of the double dynamics of this system is a type of ‘homeostasis’. Neither potentiation nor depression level becomes saturated. Both attain a stationary level determined by the coexistence of potentiation and depression in every functional synaptic population, due to the fact that a significant fraction of the neurons belong to the representation of more than one stimulus.

4.1. Robustness of the learning process

The process of synaptic structuring, together with its neural correlates, is very robust. This was corroborated by removing a number of simplifying assumptions used in the main simulation. First, we carried out simulations in which variability was introduced in the synaptic efficacies within every synaptic population (excitatory-to-inhibitory, inhibitory-to-excitatory, inhibitory-to-inhibitory, excitatory-to-excitatory potentiated and depressed). The efficacies have been drawn from a Gaussian distribution centred around the corresponding J used in the main simulation. The variance of the distribution is set to $\Delta^2 J^2$, with Δ varying between 0.1 and 0.5 (Amit & Brunel, 1997a). In none of the simulations carried out did we observe noticeable differences with respect to the phenomenology described previously.

Also, we ran simulations with randomly variable coding levels for the stimuli together with noise during stimulus presentation (Section 3.4) and we desynchronized the activation of the external selective currents to different neurons upon stimulus presentation (Section 3.5). In both cases, the process of synaptic structuring, as well as the functioning of the network as a WM, were qualitatively unaffected. Not surprisingly, large levels of noise during stimulus presentation, throughout the training, degraded the performance of the network. However, if there was no noise during training the same noise levels did not affect performance at asymptotic structuring (see Section 3.4 for details).

We are currently testing the robustness of the learning process to several additional aspects, keeping the average parameters of the network as in the main simulation.

1. Allowing variability of the ‘contrast’ factor from neuron to neuron. In each stimulation, the ‘contrast’ factors are randomly drawn from Gaussian distributions, centred on the ‘contrast’ factors in the main simulation.

2. Allowing nonselective currents during stimulation. All neurons not selective for the stimulus receive increased external currents, during presentation, with ‘contrast’ factor lower than that of the selective neurons.

3. Allowing finite excitatory and inhibitory synaptic time constants. We chose $\tau_s = 2$ ms for the fast excitatory currents and $\tau_s = 5$ ms for the inhibitory currents (see Materials and methods). We found it necessary to increase both σ_{ext}^E and J_{EI} to eliminate oscillatory behaviours (see, e.g., Traub *et al.*, 1999), to increase J_p to ensure stable, long-living delay activity, and to reduce G_{stim}^E to maintain the average steady response before structuring as in the main simulation.

Preliminary results show the phenomenology remained the same as in the main simulation.

4.2. The roles of short-term depression and NMDA

In obtaining ‘robust’ learning, a fundamental role has been played by the mechanism of short-term synaptic depression implemented (Tsodyks & Markram, 1997; Tsodyks *et al.*, 1998). It also allows for the removal of external manual intervention during training (Amit & Mongillo, 2003). Upon structuring, stimulus response tends to increase and this increase is a source of instabilities in the learning process (Del Giudice & Mattia, 2001; Amit & Mongillo, 2003). In fact, too high average emission rates upon stimulus presentation could alter the balance between excitation and inhibition, leading to the appearance of oscillatory behaviour or uncontrolled growth of global network activity. Such high rates can significantly affect the probability of synaptic transitions within the various functional synaptic populations, leading to undesirable synaptic modifications (e.g. potentiation instead of depression) which may distort the learning process. Short-term rate-dependent depression of the excitatory synaptic efficacies contains

stimulus response within functional boundaries. During stimulation the network reliably works in an asynchronous irregular spiking regime (balanced regime). This ensures that the probability of synaptic transitions is significantly different from zero, and low, only in synapses among responsive neurons (LTP) and in the synapses from responsive to nonresponsive neurons (LTD).

With short-term synaptic depression, we have found it necessary to add slow NMDA-like currents to ensure the proper functioning of selective WM activity. In its absence, after the network had structured itself properly and had stable selective delay activity states, either according to mean-field analysis (Curti *et al.*, 2004) or by direct access to these states in the simulation, the neural dynamics of the network, following the removal of a stimulus, did not reach the WM state corresponding to that stimulus. The reason is that, at the end of stimulus presentation, the fraction of available resources of the synapses corresponding to selective neurons is very low due to the relatively high emission rate (≈ 50 Hz) during stimulation. They recover on a relatively long time scale (≈ 200 – 300 ms, in simulation), and the memory of the stimulus presented may be lost. The network either relaxes to the spontaneous activity state or to a WM state corresponding to a different randomly selected stimulus.

The slow (100 ms) NMDA-like component (50%) of the excitatory recurrent current keeps track of the information of the presented stimulus while synaptic resources recover. In other words, after stimulus removal there is a selective recurrent current within the stimulated neural population. This NMDA current compensates for the temporary low availability of synaptic resources, maintaining the emission rate within the selected population at a relatively higher level. Slow NMDA-like currents also render WM activity more stable (see, e.g., Wang, 1999; Compte *et al.*, 2000; Brunel & Wang, 2001).

Such slow currents also play a role in preventing spontaneous global bursting, commonly manifested in networks of spiking neurons with short-term depressing synapses (Tsodyks *et al.*, 2000; Loebel & Tsodyks, 2002). Bursts occur because of occasional synchronous firing of a subset of excitatory cells and, with fast excitatory currents, synchronous spiking provokes an avalanche of firing activity, and nearly all neurons in the network spike within a few ms (Tsodyks *et al.*, 2000). A slower recurrent current renders inhibition more effective in controlling small spontaneous fluctuations of the activity in the excitatory population. The network operates reliably in an asynchronous firing regime. At the same time, the network maintains its ability to generate fast-developing activity in response to stimulus presentation. This bursting regime is, however, limited to neurons belonging to a selective population, which consists of a small fraction of the entire excitatory population.

4.3. LTP–LTD transition probabilities and neural activity

The average transition probabilities per presentation increase with training, as a consequence of the increase with structuring of the average emission rates upon stimulation. The shorter the average ISI, the larger the transition probabilities. In particular, the initial burst can significantly affect the transition probabilities. During the initial transient, due to the temporary imbalance between excitation and inhibition and to the large availability of synaptic resources, single neurons can fire multiple action potentials with short ISIs. In fact, reducing the amplitude of the peak response (and its relative variation through training) results in the reduction of the transition probabilities and of its increase during training. (Section 3.5).

It must be noted, however, that the relative variation of the LTP–LTD transition probabilities in the course of training is not fully accounted for by the corresponding variation of the average ISI upon

stimulation. This is due to the low CVs during stimulus presentation (Section 3.2.5). For the same average ISI (i.e., the same mean emission rate), a larger CV implies a long tail of the underlying distribution of firing times. This reduces the probability of a synaptic transition. One also expects that higher CVs upon stimulation may produce low and quite constant probabilities of synaptic transitions throughout training. Larger CVs can be obtained by reducing the average emission rates upon stimulus presentation, by increasing the level of external noise during stimulation and by allowing Hebbian plasticity in the inhibitory-to-excitatory synapses.

From the theoretical point of view, obtaining transition probabilities low and constant through the training history has the important consequence that the learning theory of Brunel *et al.* (1998) becomes applicable together with its relative, the mean-field description of the stationary states of the neural system (Curti *et al.*, 2004). Without the constancy of the synaptic transition probabilities the first step is not possible. The second can still be partially saved, with much reduced effectiveness, by running the simulation up to a given stage and then using the resulting potentiation levels in different functional synaptic populations to set up an instantaneous mean-field analysis. On the other hand, the rise in probabilities and their dependence on the burst amplitude may also have a positive side in the development of associations between delay activity distributions in associative learning, as in the pair-associate paradigm or in learning sequences of stimuli with a fixed order. Study of the related effects is under way.

4.4. Achievements and perspectives

The model network captures several important aspects of the experimentally observed phenomenology. (i) Both excitatory and inhibitory neurons respond selectively to the stimuli; they respond with the same coding level and with roughly equal mean emission rates (Tamura *et al.*, 2004). (ii) The discharge patterns of the neurons are quite irregular, and the distribution of the firing times is characterized by a long tail (Section 3.2.5). (iii) The distributions of the selective average emission rates, during stimulus presentation as well as during the subsequent delay interval, are wide and similar to those experimentally reported (Section 3.2.5). (iv) The time course of the stimulus response is consistent with profiles observed in *in vivo* recordings (Tamura & Tanaka, 2001; Tamura *et al.*, 2004): for excitatory neurons, fast initial transient at a high rate followed by a steady response at a lower rate; for inhibitory neurons, tonic response throughout the stimulation.

Because of its biological credibility, the model may constitute a useful tool in tracing learning-related modifications of neural activity in experiments, as well as in designing new, informative experiments. It also makes experimentally testable statements. The model predicts an increase in the stimulus response with training. Presentation of familiar stimuli produces higher rates than those produced by novel stimuli. At present there is barely preliminary evidence for this. In Messinger *et al.* (2001) and Holscher *et al.* (2004), the emission rate of single neurons in perirhinal cortex and area TE in inferotemporal cortex (IT) during training for a delay task increases significantly with the number of presentations per stimulus.

Another neural correlate of learning is the narrowing of the single-cell tuning curves, i.e. responses to the stimuli which are poor before learning are completely suppressed, or reduced, with training (see, e.g., Rainer & Miller, 2000; Zohary *et al.*, 1994; Holscher *et al.*, 2004). Rainer & Miller (2000) also report that the selectivity to the familiar stimuli is more robust with respect to the stimulus degradation (see also Amit *et al.*, 1997). Both effects are easily accounted for in the

model. As the stimulus becomes familiar (i.e. it is repeatedly presented), the fraction of potentiated synapses among the neurons selective to it increases. In this way, even when the stimulus is degraded by noise (see Section 3.4), the nonstimulated selective cells emit at enhanced rates because of the strong recurrent synaptic efficacies. Similarly, with training, the fraction of potentiated synapses from selective to nonselective neurons decreases. Consequently, upon presentation of a given stimulus the excitatory currents afferent on the nonselective neurons decrease, reducing their emission rates. This produces the narrowing of the tuning curves.

In the simulations reported we did not observe this effect directly because during stimulus presentation nonselective cells are practically quiescent, throughout the training. However, we did observe a decrease in the average depolarization level of the nonselective cells with training, indicating a corresponding reduction in the afferent current. In preliminary simulations we increased the nonselective afferent current during stimulation. This, on the one hand, may be interpreted as a change in the subject's attention in viewing the stimulus; on the other hand it endows nonselective cells with significantly higher emission rates without harming the structuring process (see Section 3.4). In these simulations we observed the narrowing of the single-cell tuning curves as training proceeded.

In this network, one can naturally study the dependence of the structuring and of the associated neural dynamics on the presentation protocol. The present study reports a basic case in which the set of stimuli is stationary (i.e. no stimuli are added or removed from the training set), and stimuli are presented in a random, uncorrelated way at the same rate. Moreover, the trial is an elementary stimulus–delay pair. The simulations can be applied to modelling the neural correlates observed in more elaborate tasks, starting from the simple delay-match-to-sample with fixed-order sequences (Miyashita, 1988), the pair-associate matching (see, e.g., Erickson & Desimone, 1999; Messinger *et al.*, 2001; Naya *et al.*, 2003), and task switching during delay (Naya *et al.*, 1996). Several of these tasks have been modelled with promising results (Amit *et al.*, 1994; Brunel, 1994; Brunel, 2003; Mongillo *et al.*, 2003), but not yet with a fully embedded microscopic synapse. Such applications may reinforce the credibility of the modelling elements, and may also hide some surprises and lead to new predictions.

Adding or removing stimuli from the training set, or varying their presentation rates, could expose learning-related modifications in experiments in a more effective way. Finally, the presentation protocol could be relevant with respect to the issues such as learning and forgetting rates (Brunel *et al.*, 1998), and the storage capacity (Curti *et al.*, 2004). Preliminary simulations show that, in some cases, the network is unable to store a given number stimuli if they are all presented in the same training session. By contrast, if training is made first on a subset of the stimuli and then on the entire set, the network develops selective delay activity for all stimuli. Such a feature of training could account for a common strategy for memorising a long text, as for instance a poem.

Last, among many possible additional ones, we mention the prospect of using such a network to investigate the unsupervised development of neural representations expressed by selective delay activity states. The network could be trained on an arbitrary set of stimuli, not necessarily random or of fixed coding level; they could, for example, be pixelised images of real objects. The neural dynamics together with the synaptic plasticity would create delay activity representations for the stimuli which are likely to be quite different from the presented stimuli and would probably give rise to higher selectivity (fewer correlations, fewer overlaps) and lower coding. Such an outcome is imaginary at this stage but, if verified, could lead to the

resolution of the puzzle of the collapse of the delay-activity representations of many images into very restricted columns.

Many of these issues are currently under investigation.

Acknowledgements

This study was supported by the Center of Excellence Grant 'Changing Your Mind' of the Israel Science Foundation and the Center of Excellence Grant 'Statistical Mechanics and Complexity' of the INFN, Roma-I. We would like to thank A. Bernacchia for helpful discussions and assistance in data analysis. G.M. would like to thank G. La Camera for helpful discussions and useful comments on an early version of the manuscript.

Abbreviations

μ , no. of a given stimulus in a set of stimuli; CV, coefficient of variation; ISI, interspike interval; LTD, long-term depression; LTP, long-term potentiation; v_{peak} , peak response (maximal population activity level in 10-ms bins during stimulus presentation); v_{steady} , steady response (average activity level in the last 350 ms of stimulus presentation); pres/st, presentations per stimulus; sel \rightarrow nonsel, selective \rightarrow nonselective, i.e. pre- and postsynaptic neurons responding to different stimuli; sel \rightarrow sel, selective \rightarrow selective, i.e. pre- and postsynaptic neurons responding to the same stimulus; $_{\text{ns}}$, selective \rightarrow nonselective; $_{\text{ss}}$, selective \rightarrow selective; WM, working memory; X , internal analogue variable.

References

- Amit, D.J. (1998) Simulation in neurobiology — theory or experiment? *TINS*, **21**, 231–237.
- Amit, D.J., Bernacchia, A. & Yakovlev, V. (2003) Multiple-object working memory – A model for behavioral performance. *Cereb. Cortex*, **13**, 435–443.
- Amit, D.J. & Brunel, N. (1997a) Dynamics of recurrent network of spiking neurons before and following learning. *Network*, **8**, 373–404.
- Amit, D.J. & Brunel, N. (1997b) Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cereb. Cortex*, **7**, 237–252.
- Amit, D.J., Brunel, N. & Tsodyks, M.V. (1994) Correlations of cortical hebbian reverberations: theory versus experiment. *J. Neurosci.*, **14**, 6435–6445.
- Amit, D.J., Fusi, S. & Yakovlev, V. (1997) A paradigmatic working memory (attractor) cell in cortex. *Neural Computation*, **9**, 1071–1092.
- Amit, D.J. & Mongillo, G. (2003) Spike-driven synaptic dynamics generating working memory states. *Neural Computation*, **15**, 565–596.
- Brunel, N. (1994) Dynamics of an attractor neural network converting temporal into spatial correlations. *Network: Computation Neural Systems*, **5**, 449–470.
- Brunel, N. (2003) Dynamics and plasticity of stimulus-selective persistent activity in cortical network models. *Cerebral Cortex*, **13**, 1151–1161.
- Brunel, N., Carusi, F. & Fusi, S. (1998) Slow stochastic hebbian learning of classes of stimuli in a recurrent neural network. *Network*, **9**, 123–152.
- Brunel, N. & Wang, X.-J. (2001) Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition. *J. Comput. Neurosci.*, **11**, 63–85.
- Compte, A., Brunel, N., Goldman, R., Rieke, P.S. & Wang, X.-J. (2000) Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cereb. Cortex*, **10**, 910–923.
- Curti, E., Mongillo, G., La Camera, G. & Amit, D.J. (2004) Mean-field and capacity in realistic networks of spiking neurons storing sparsely coded random memories. *Neural Computation*, **16**, 2597–2637.
- Del Giudice, P., Fusi, S., Badoni, D., Dante, V. & Amit, D.J. (1998) Learning attractors in an asynchronous, stochastic electronic neural network. *Network*, **9**, 183–205.
- Del Giudice, P. & Mattia, M. (2001) Long and short-term synaptic plasticity and the formation of working memory: a case study. *Neurocomputing*, **38–40**, 1175–1180.
- Erickson, C.A. & Desimone, R. (1999) Responses of macaque perirhinal neurons during and after visual stimulus association learning. *J. Neurosci.*, **19**, 10404–10416.
- Fusi, S., Annunziato, M., Badoni, D., Salamon, A. & Amit, D.J. (2000) Spike-driven synaptic plasticity: theory, simulation, VLSI implementation. *Neural Computation*, **12**, 2227–2258.
- Fusi, S. & Mattia, M. (1999) Collective behavior of networks with linear (vlsi) integrate and fire neurons. *Neural Computation*, **11**, 633–652.
- Fuster, J.M. & Alexander, G.E. (1971) Neuron activity related to short-term memory. *Science*, **173**, 652–654.
- Holscher, C., Rolls, E.T. & Xiang, J. (2004) Perirhinal cortex neuronal activity related to long-term familiarity memory in the macaque. *Eur. J. Neurosci.*, **18**, 2037–2046.
- Loebel, A. & Tsodyks, M.V. (2002) Computation by ensemble synchronization in recurrent networks with synaptic depression. *J. Comput. Neurosci.*, **13**, 111–124.
- Mattia, M. & Del Giudice, P. (2000) Efficient event-driven simulation of large networks of spiking neurons and dynamical synapses. *Neural Computation*, **12**, 2305–2329.
- Messinger, A., Squire, L.R., Zola, S.M. & Albright, T.D. (2001) Neuronal representations of stimulus associations develop in the temporal lobe during learning. *PNAS*, **98**, 12239–12244.
- Meyer, P.L. (1965) *Introductory Probability and Statistical Applications*. Reading, MA: Addison-Wesley.
- Miyashita, Y. (1988) Neural correlate of visual associative long-term memory in the primate temporal cortex. *Nature*, **335**, 817–820.
- Miyashita, Y. & Chang, H.S. (1988) Neuronal correlate of pictorial short-term memory in the primate temporal cortex. *Nature*, **331**, 68–70.
- Mongillo, G., Amit, D.J. & Brunel, N. (2003) Retrospective and prospective persistent activity induced by hebbian learning in a recurrent cortical network. *Eur. J. Neurosci.*, **18**, 2011–2024.
- Nakamura, K. & Kubota, K. (1995) Mnemonic firing of neurons in the monkey temporal pole during a visual recognition memory task. *J. Neurophysiol.*, **74**, 162–178.
- Naya, Y., Yoshida, M. & Miyashita, Y. (2003) Forward processing of long-term associative memory in monkey inferotemporal cortex. *J. Neurosci.*, **23**, 2861–2871.
- Naya, Y., Sakai, K. & Miyashita, Y. (1996) Activity of primate inferotemporal neurons related to a sought target in pair association task. *PNAS*, **93**, 2664–2669.
- Rainer, G. & Miller, E.K. (2000) Effects of visual experience on the representation of objects in the prefrontal cortex. *Neuron*, **27**, 179–189.
- Rauch, A., La Camera, G., Luscher, H.-R., Senn, W. & Fusi, S. (2003) Neocortical pyramidal cells respond as integrate-and-fire neurons to in vivo-like input currents. *J. Neurophysiol.*, **90**, 1598–1612.
- Shadlen, M.N. & Newsome, W.T. (1998) The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J. Neurosci.*, **18**, 3870–3896.
- Sofky, W. & Koch, C. (1993) The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J. Neurosci.*, **13**, 334–350.
- Tamura, H., Kaneko, H., Kawasaki, K. & Fujita, I. (2004) Presumed inhibitory neurons in the macaque inferior temporal cortex: visual response properties and functional interactions with adjacent neurons. *J. Neurophysiol.*, **91**, 2782–2796.
- Tamura, H. & Tanaka, K. (2001) Visual response properties of cells in the ventral and dorsal parts of the macaque inferotemporal cortex. *Cerebral Cortex*, **11**, 384–399.
- Traub, R.D., Jefferys, J.G.R. & Ni, Whittington, A. (1999) *Fast Oscillations in Cortical Circuits*. The MIT Press, Cambridge, MA.
- Tsodyks, M.V. & Markram, H. (1997) The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *PNAS*, **94**, 719–723.
- Tsodyks, M.V., Pawelzik, K. & Markram, H. (1998) Neural networks with dynamic synapses. *Neural Computation*, **10**, 821–835.
- Tsodyks, M. & Uziel, A., L. & Markram, H. (2000) Synchrony generation in recurrent networks with frequency-dependent synapses. *J. Neurosci.*, **20** (RC50), 1–5.
- Wang, X.-J. (1999) Synaptic basis of cortical persistent activity: the importance of NMDA receptors to working memory. *J. Neurosci.*, **19**, 9587–9603.
- Zohary, E., Celebrini, S., Britten, K. & Newsome, W. (1994) Neuronal plasticity that underlies improvement in perceptual performance. *Science*, **265**, 1289–1292.