

Oscillations, Phase-of-Firing Coding, and Spike Timing-Dependent Plasticity: An Efficient Learning Scheme

Timothée Masquelier,¹ Etienne Hugues,¹ Gustavo Deco,¹ and Simon J. Thorpe²

¹Theoretical and Computational Neuroscience Group, Center for Brain and Cognition, Universitat Pompeu Fabra, Barcelona 08018, Spain, ²Centre de Recherche Cerveau et Cognition, Université Toulouse 3, Centre National de la Recherche Scientifique, Faculté de Médecine de Rangueil, Toulouse 31062, France

Recent experiments have established that information can be encoded in the spike times of neurons relative to the phase of a background oscillation in the local field potential—a phenomenon referred to as “phase-of-firing coding” (PoFC). These firing phase preferences could result from combining an oscillation in the input current with a stimulus-dependent static component that would produce the variations in preferred phase, but it remains unclear whether these phases are an epiphenomenon or really affect neuronal interactions—only then could they have a functional role. Here we show that PoFC has a major impact on downstream learning and decoding with the now well established spike timing-dependent plasticity (STDP). To be precise, we demonstrate with simulations how a single neuron equipped with STDP robustly detects a pattern of input currents automatically encoded in the phases of a subset of its afferents, and repeating at random intervals. Remarkably, learning is possible even when only a small fraction of the afferents (~10%) exhibits PoFC. The ability of STDP to detect repeating patterns had been noted before in continuous activity, but it turns out that oscillations greatly facilitate learning. A benchmark with more conventional rate-based codes demonstrates the superiority of oscillations and PoFC for both STDP-based learning and the speed of decoding: the oscillation partially formats the input spike times, so that they mainly depend on the current input currents, and can be efficiently learned by STDP and then recognized in just one oscillation cycle. This suggests a major functional role for oscillatory brain activity that has been widely reported experimentally.

Introduction

Whether spike times contain additional information with respect to time-averaged firing rates—a theory referred to as “temporal coding”—is still a matter of debate. It has been argued that the speed of processing in sensory pathways makes it difficult to obtain a reliable measure of the firing rates of individual neurons within the available time window (Thorpe and Imbert, 1989; Gautrais and Thorpe, 1998). However, the firing rate of a neuron population with similar selectivity could be reliably estimated over an arbitrary small time window provided the population is large enough. This concept, referred to as “population coding” (Pouget et al., 2000), has often been proposed to explain how the brain might cancel out the variability frequently observed in individual spike trains (Softky and Koch, 1993; Shadlen and Newsome, 1998).

However, a recent experimental study effectively ruled out the possibility that only population coding is used, at least in the mouse retina, simply because its bandwidth is too low to account for the observed behavioral performance (Jacobs et al., 2009). It

appears that some information has to be encoded in the spike times. This raises a number of important questions, such as first, how (somewhat) reliable spike times can be produced; second, how those spike times can be decoded, that is, how a neuron response can be made selective to a given input spike pattern; and third, how this selectivity could be learned. So far, most of the research effort has focused on the first question.

In some experimental paradigms, a given sensory system is in a resting state before being presented with a stimulus. The stimulus onset then provides an external reference time, and it has been repeatedly demonstrated that the response latencies with respect to it can potentially encode information (Celebrini et al., 1993; Gawne et al., 1996; Albrecht et al., 2002; Kiani et al., 2005). The relative latencies are sometimes even more informative (Johansson and Birznieks, 2004; Chase and Young, 2007; Gollisch and Meister, 2008). Benchmark 3 (see below) is an attempt to model how selectivity can emerge in those paradigms.

However most of the usual cognitive processes are ongoing and cannot make use of an external event as a reference. In these cases, it has been suggested that when the local field potential presents a prominent oscillation, it may serve as a time reference, and a number of mechanisms have been proposed to account for stimulus-dependent phase locking (Buzsáki and Chrobak, 1995; Hopfield, 1995; Mehta et al., 2002; Brody and Hopfield, 2003; Buzsáki and Draguhn, 2004; Lisman, 2005; Fries et al., 2007) and have experimental support *in vitro* (McLelland and Paulsen, 2009). Consistent with those models, more and more *in vivo* evidence has been found for the so called “phase-of-firing coding” (PoFC) (see Table 1).

Received May 11, 2009; revised Aug. 19, 2009; accepted Sept. 3, 2009.

This research was supported by the Fyssen Foundation, the Centre National de la Recherche Scientifique, and Specific Targeted Research Project Decisions-in-Motion (IST-027198). We thank Dan Goodman and Romain Brette for having developed the Brian simulator, both efficient and user-friendly, and for the quality of their support.

Correspondence should be addressed to Timothée Masquelier, Theoretical and Computational Neuroscience Group, Center for Brain and Cognition, Universitat Pompeu Fabra, C/Roc Boronat, 138, 08018 Barcelona, Spain. E-mail: timothee.masquelier@alum.mit.edu.

DOI:10.1523/JNEUROSCI.2207-09.2009

Copyright © 2009 Society for Neuroscience 0270-6474/09/2913484-10\$15.00/0

Table 1. Experimental evidence for phase-of-firing coding

Animal	Recording site	Frequency band	What is coded	Reference
Locust	Mushroom body	Beta–gamma (20–30 Hz)	Odor identity	Perez-Orive et al. (2002); Cassenaer and Laurent (2007)
Honeybee	Antennal lobe	Gamma	Odor identity	Stopfer et al. (1997)
Zebrafish	Olfactory bulb	Gamma	Odor identity	Friedrich et al. (2004)
Rat	Olfactory bulb	Gamma	Odor identity	Eeckman and Freeman (1990); David et al. (2009)
Rabbit	Olfactory bulb	Gamma	Odor identity	Kashiwadani et al. (1999)
Rat	Hippocampus (CA1 and CA3)	Theta	Spatial location	O'Keefe and Recce (1993); Mehta et al. (2002)
Rat	Entorhinal cortex	Theta	Spatial location	Hafting et al. (2008)
Cat	V1	Gamma	Line orientation	König et al. (1995); Fries et al. (2001)
Macaque	V1	Delta	Visual (motion) features	Montemurro et al. (2008)
Macaque	V4	Theta	Memorized image	Lee et al. (2005)
Macaque	Auditory cortex	Theta	Auditory stimulus identity	Kayser et al. (2009)

Despite this converging evidence, the second and third questions remain largely unanswered. Specifically, we still have little idea how might the brain decode information contained in phase information, nor is clear how those patterns could be learnt. Here we demonstrate through simulations that the now well established spike timing-dependent plasticity (STDP) (for a recent review, see Caporale and Dan, 2008) can efficiently solve the problem (benchmark 4 below). The ability of STDP to learn and detect spike patterns had been noted before in continuous activity (Masquelier et al., 2008), but not in an oscillatory mode like here, and it turns out that learning is significantly facilitated in this case.

Materials and Methods

The problem: detecting a repeating activation pattern. We want to solve a difficult abstract problem with which a single neuron could be confronted: detecting the existence of patterns of input activation involving an unknown subset of the afferents, under conditions where the duration of the patterns is unpredictable and where the patterns recur at unpredictable intervals.

To be precise, let us consider a matrix of real values, corresponding to for now abstract neuronal activation levels, scaled in $[0, 1]$ (see Fig. 1). The number of rows n is the number of neurons in the population (here 2000). Each column corresponds to the neuronal activation pattern for a certain time Δt , drawn from an exponential distribution with mean 250 ms. Then the neurons' activation levels all change at the same time, adopting the values of the next column and so on. Now consider the situation where a fraction x (10% in the baseline simulations) of a given column, referred to as “the pattern,” is copied at random intervals (so that the interpattern interval has an exponential distribution with mean 1250 ms—the pattern is thus there 20% of the time, and “concerns” only x percent of the population). The copied values are referred to as pattern activation levels. Note that in the figure, and in this paper in general, we have arranged it so that the neurons in the pattern have the lowest indexes to make the pattern easier to see, but of course the indexes are arbitrary.

Here we investigated various biologically plausible mechanisms to convert the abstract activation levels into spikes, and tested whether a downstream neuron equipped with STDP was able to detect and learn the repeating pattern. To make the problem more difficult, we normalized the lines and columns of the matrix, such that the time-averaged activation levels are uniform over the population, and the population-averaged activation levels are constant along the whole simulation. We built such matrices through successive normalizations, and we would be happy to make them available to anyone wanting to benchmark other encoding/decoding algorithms.

Neuronal models. The simulations were run with Brian (<http://www.briansimulator.org/>), a new Python-based clock-driven spiking neural network simulator (Goodman and Brette, 2008). The code has been made available on ModelDB (<http://senselab.med.yale.edu/ShowModel.asp?model=123928>).

In benchmark 1, we used Poisson input neurons. The matrix activation levels of Figure 1 were linearly mapped to their firing rates f . At each

time step $dt = 0.1$ ms, the probability of emitting a spike was proportional to $f \cdot dt$.

In all other simulations we used a leaky integrate-and-fire (LIF) neuron model, receiving a current made of a time-dependent input $I(t)$ plus a Gaussian white noise $\xi(t)$ [with $\langle \xi(t) \rangle = 0$ and $\langle \xi(t)\xi(s) \rangle = \delta(t - s)$]. The membrane potential obeys to the following Langevin equation:

$$\tau_m \frac{dV}{dt} = -V + E_1 + RI(t) + \sigma \cdot \xi \cdot \sqrt{\tau_m}, \quad (1)$$

where V is the membrane potential, $E_1 = -70$ mV is its resting value, $\tau_m = 20$ ms is the membrane time constant, $R = 10$ M Ω is its resistance, and $\sigma = 0.015 \cdot (V_t - V_r) = 0.09$ mV is the standard deviation of the noise. Whenever the threshold $V_t = -54$ mV is reached, a postsynaptic spike is emitted, and the membrane potential is reset to $V_r = -60$ mV and clamped there for a refractory period of 1 ms. The differential equation is solved numerically (Euler method) with a time step of $dt = 0.1$ ms.

We used $n = 2000$ input neurons, which receive input currents I_1, \dots, I_n that are affinely mapped on the matrix activation levels of Figure 1. In benchmark 4, they additionally receive a common sinusoidal drive:

$$i(t) = \frac{a}{2} \cdot \sin(2\pi f \cdot t - \pi), \quad (2)$$

where $a = 0.15 \cdot I_{\text{thr}} = 0.24$ nA [$I_{\text{thr}} = (V_t - E_1)/R$ is the threshold current] is the peak-to-peak-amplitude of the oscillation and $f = 8$ Hz its frequency.

As can be seen in Figure 2, one downstream neuron is connected to all the input neurons with normalized synaptic weights $w_1(t), \dots, w_n(t)$ in $[0, 1]$, so that its input current in Equation 1 is as follows:

$$I(t) = I_{\text{max}} \sum_{j=1}^n \sum_{k=1}^{n_j} w_j(t_j^k) \cdot \exp\left(-\frac{t - t_j^k}{\tau_s}\right) \cdot \Theta(t - t_j^k), \quad (3)$$

where $t_j^1 \dots t_j^{n_j}$ denote the times of the spikes emitted by input neuron j , $\tau_s = 5$ ms is the synapse time constant, Θ is the Heaviside step function:

$$\Theta(s) = \begin{cases} 1 & \text{if } s \geq 0 \\ 0 & \text{if } s < 0 \end{cases}, \quad (4)$$

and I_{max} corresponds to the input current contributed by one input spike right after it is received through a maximally reinforced synapse ($w = 1$). We used various values for I_{max} (see below). The initial synaptic weights $w_1(0), \dots, w_n(0)$ were randomly picked from a uniform distribution between 0 and an upper bound chosen such that $\bar{w} \cdot I_{\text{max}} = 8.6 \times 10^{-12}$ A, where \bar{w} is the average of $w_1(0), \dots, w_n(0)$. This means that at the beginning of the simulations, the equivalent of 186 synchronous spikes (that is, a synchronized activity of about 10% of the network) is needed to reach the threshold from the resting state.

STDP model. Unless said otherwise we used a classic “all-to-all” additive weight update rule (Kempster et al., 1999; Song et al., 2000), where all pairs (i, j) of postsynaptic and presynaptic spikes contribute to the weight change:

$$\Delta w_j = \begin{cases} a^+ \cdot \exp\left(\frac{t_j - t_i}{\tau^+}\right) & \text{if } t_j \leq t_i \text{ (LTP, applied at time } t_i) \\ -a^- \cdot \exp\left(-\frac{t_j - t_i}{\tau^-}\right) & \text{if } t_j > t_i \text{ (LTD, applied at time } t_j) \end{cases} \quad (5)$$

In one simulation of benchmark 4, we also implemented the “nearest spike” mode in which weight updates caused by one postsynaptic spike are restricted at each synapse to the latest preceding presynaptic spikes and the earliest following one (Burkitt et al., 2004).

We used $\tau^+ = 16.8$ ms and $\tau^- = 33.7$ ms, as observed experimentally (Bi and Poo, 2001). Since the additive STDP is not naturally bounded we clipped the weights to $[0, 1]$ after each weight update. We used $a^+ = 0.005$, and various ratios a^-/a^+ (see below).

Mutual information. We used information theory to quantify how good the postsynaptic neuron is at detecting the pattern after convergence. To be specific, we discretized the $[800$ s, 1000 s] period into 125 ms time bins. For each of those bins, we determined whether the stimulus (or pattern) was present more than half of the time (case referred to as s) or not (\bar{s}), and whether the postsynaptic neuron emitted at least one spike (r) or not (\bar{r}). The mutual information between the postsynaptic response and the presence of the stimulus is then given by the following:

$$MI = P(r, s) \log_2 \left(\frac{P(r, s)}{P(r) \cdot P(s)} \right) + P(\bar{r}, s) \log_2 \left(\frac{P(\bar{r}, s)}{P(\bar{r}) \cdot P(s)} \right) + P(r, \bar{s}) \log_2 \left(\frac{P(r, \bar{s})}{P(r) \cdot P(\bar{s})} \right) + P(\bar{r}, \bar{s}) \log_2 \left(\frac{P(\bar{r}, \bar{s})}{P(\bar{r}) \cdot P(\bar{s})} \right). \quad (6)$$

Note that in signal detection terms, the first term corresponds to “hits,” the second to “misses,” the third to “false alarms,” and the last one to “correct rejections.”

A perfect detector would lead to $P(r, \bar{s}) = P(\bar{r}, s) = 0$, $P(r, s) = P(s) = P(r)$, and $P(\bar{r}, \bar{s}) = P(\bar{s}) = P(\bar{r})$. Therefore, an upper bound on the mutual information is the following:

$$MI_{\max} = -P(s) \log_2 P(s) - P(\bar{s}) \log_2 P(\bar{s}) = H_s, \quad (7)$$

which is the entropy of the stimulus. Here $P(s) = 1/5$, leading to $MI_{\max} \approx 0.72$ bits.

Results

We benchmarked several mechanisms to encode the matrix activation levels of Figure 1 into spikes. The criterion was the degree to which a downstream neuron equipped with STDP (Fig. 2) was able to detect and learn the repeating pattern, starting from uniformly distributed random synaptic weights.

Table 2 summarizes the different benchmarks. We used Poisson neurons in benchmark 1, and LIF neurons in the rest of them (Eq. 1). In benchmark 3, we sometimes reset the potential of all the afferents simultaneously, and in benchmark 4 we use an additional oscillatory drive (Eq. 2). Each benchmark is detailed below.

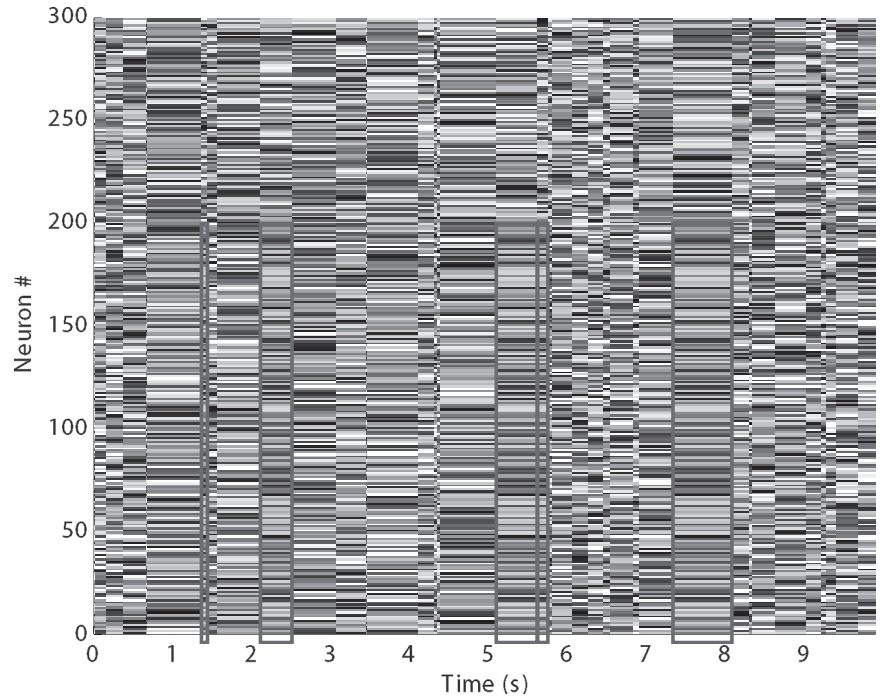


Figure 1. Overview of the problem. Here we plot the abstract activation levels scaled in $[0, 1]$ (represented by the gray levels, white corresponding to 1 and black to 0) of a group of neurons (y -axis), as a function of time (x -axis). For space reasons we have only represented the first 300 of the 2000 afferents used in the simulations. All the activation levels change simultaneously every Δt , drawn from an exponential distribution with mean 250 ms. At random intervals, drawn from an exponential distribution with mean 1250 ms, the pattern of activation levels for a subgroup of neurons (here indexed $0 \dots 199$) repeats (rectangles). These repeating values are referred to as pattern activation levels. Our goal is to find an encoding mechanism such that a downstream neuron equipped with STDP could detect and become selective to this repeating pattern.




Figure 2. Network architecture. Afferents $1 \dots n$ are shown on the left. They receive static input currents $I_1 \dots I_n$. In benchmark 4, they additionally receive a common oscillatory drive $i(t)$. They are all connected, through excitatory synapses with weights $w_1 \dots w_n$, to one downstream neuron equipped with STDP.

Benchmark 1: Poisson input neurons

Poisson neurons are extensively used among the computational neuroscience community. The first reason is that it is commonly believed that the variability of neuronal responses is well described by Poisson statistics (Softky and Koch, 1993; Shadlen and Newsome, 1998). Note that this view is under challenge because a number of recent studies have shown that neuronal responses are too reliable for the Poisson hypothesis to be tenable in the lateral geniculate nucleus (Liu et al., 2001), the retina (Uzzell and Chichilnisky, 2004), and the inferotemporal cortex (Amarasingham et al., 2006). The second reason for the widespread use of Poisson-based modeling is that they lead to mathematically tractable problems.

It thus seemed natural to include a Poisson coding scheme in our benchmark. The matrix activation levels were linearly

Table 2. Summary of the different benchmarks

	Neuronal model	Activation levels mapped to	Resets	Additional oscillatory drive
1	Poisson	Firing rate	No	No
2	LIF	Static input current	No	No
3	LIF	Static input current	Yes	No
4	LIF	Static input current	No	Yes

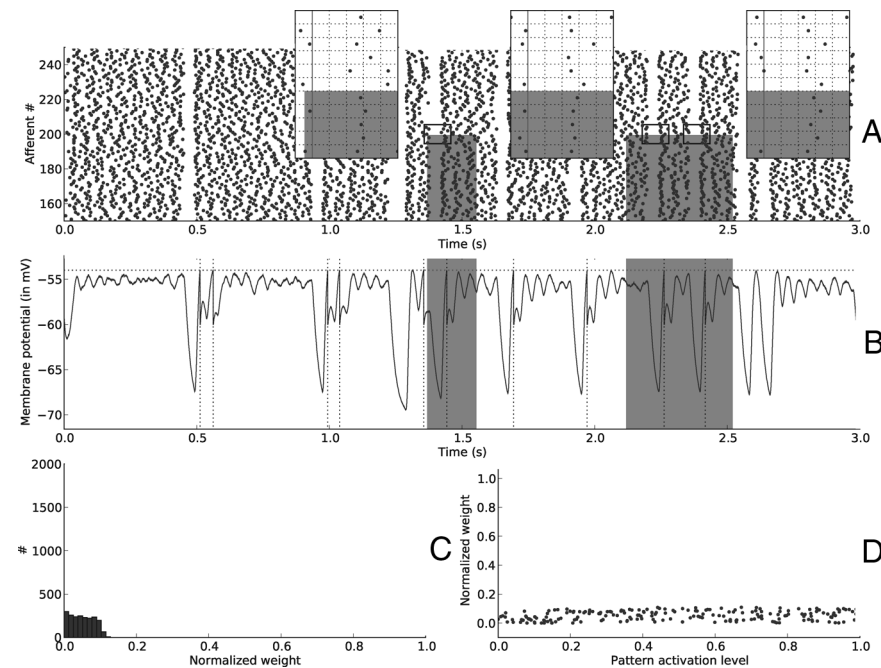


Figure 3. Resets—beginning of learning. **A**, Input spike trains from afferents 150...250, on the [0, 3 s] period. Note the silent period lasting ~50 ms after each reset. Gray rectangles designate the periods when the pattern is present, and the afferents that are involved in it (here 0...199). Three insets zoom on adequate periods to illustrate that afferents involved in the pattern (0...199) have first spike times after resets (indicated by vertical lines) that are the same (except for the noise) for different pattern presentations (the grid has a 20 ms time step and is aligned with the reset, marked with a solid line). This is not true for afferents that are not involved in the pattern (200...1999) (however, those afferents have the same latencies in the two right insets because their activation levels did not change), nor for resets outside pattern periods. It is this repeating “first spike wave” that STDP can detect and learn. **B**, Postsynaptic membrane potential as a function of time. Notice how it drops after each reset, because no input spikes are received. This is followed by a period when the membrane potential oscillates, because input spikes tend to be synchronized after a reset. The horizontal dotted line shows the threshold. A postsynaptic spike, indicated by a vertical dotted line, is emitted whenever it is reached. Again gray rectangles indicate pattern periods but not surprisingly, the response is not selective to them at this stage. **C**, Distribution of synaptic weights at $t = 3$ s. Weights are still almost uniformly distributed at this stage because few STDP updates have been made. **D**, Synaptic weights for afferents 0...199 (involved in the pattern) as a function of their corresponding pattern activation levels. Both variables are uncorrelated at this stage.

mapped to the firing rates of the input neurons. A downstream neuron equipped with STDP integrated these non-homogeneous Poisson spike trains. An exhaustive search over the range of firing rates used, the initial weight values, and the parameters I_{\max} and a^-/a^+ failed to reveal any situation where the neuron was able to learn the repeating pattern. Usually, the downstream neuron either became silent after too much long-term depression (LTD), or instead started bursting after all the synapses had been maximally reinforced. With a carefully tuned a^-/a^+ ratio, it was also possible to reach a balanced regime like that in the study by Song et al. (2000), but the discharges were not related to the presence of the pattern.

This failure is not surprising. STDP is sensitive to correlations in the input spike times (Kempster et al., 1999; Song et al., 2000; Güttig et al., 2003), not in the rates, like we have here. Detecting such rate correlations would be possible with a rate-based Hebbian rule, but estimating the rates takes time: it would require at

least a few interspike intervals (ISIs) (Gautrais and Thorpe, 1998). Thus the pattern would need to be present for at least a few typical ISIs to be detected. We will see how global resets or oscillations provide a way to detect the repeated pattern in just one ISI.

Benchmark 2: LIF input neurons

Another natural idea is to consider that the matrix activation levels are currents, and to feed them into LIF neurons. This corresponds to Figure 2 with $i(t) = 0$. Feeding a constant suprathreshold current into a LIF neuron like here leads to constant ISIs (i.e., periodic firing), except for the noise, and the relation between the input current and the ISI is bijective. This means that from one pattern presentation to another the ISI for each input neuron is the same (and thus so is the rate). However, because the neurons can be in a different state (i.e., have a different membrane potential) when the pattern is presented, depending on their history, the offsets of their periodic spike trains will differ. Because of these shifts, spike times are not the same between pattern presentations, and STDP fails again to detect them, despite an exhaustive search on the parameters I_{\max} and a^-/a^+ and the range of input currents.

Benchmark 3: LIF input neurons

+ resets

To eliminate the impact of the neuron history and align the spike trains, one can think of a “global reset” operation, that is, every now and then resetting the potential of all the LIF input neurons to their reset values. In the brain, this may be what is happening in discrete sensory processing such as saccades in vision, or sniffs in olfaction (Uchida et al., 2006). Humans perform 3–5 saccades per second; we thus chose to implement a reset every Δt , drawn from a normal distribution with mean 250 ms and standard deviation 125 ms. Note that what is happening after a

reset is also equivalent to what would be happening after presenting a stimulus to a system originally in a resting state. Thus this benchmark also corresponds to the “stimulus onset paradigms” mentioned in Introduction.

We first set the proportion of afferents involved in the pattern to be $x = 10\%$. The matrix activation levels were then affinely mapped into input currents between $1.0 \cdot I_{\text{thr}}$ and $1.05 \cdot I_{\text{thr}}$, which led to a mean input spike rate of 15.6 Hz. We chose $I_{\max} = 0.16$ nA (meaning that the equivalent of 10 synchronous spikes arriving through maximally reinforced synapses are needed to reach the threshold from the resting state) and $a^-/a^+ = 0.78$, after an exhaustive 2D grid search (geometric progressions with ratios $1.05^2 \approx 1.10$ and 1.05, respectively) to maximize the final mutual information as defined above.

Figure 3 illustrates the situation at the beginning of the simulation. Figure 3A plots the spike trains from a subset of the afferents, during the first 3 s of simulation. Three zoom insets

illustrate a key phenomenon: each time a reset occurs while the pattern is present, the first spikes after the reset for the afferents in the pattern have the same latencies, because these latencies only depend on the pattern activation levels (except for the noise, here responsible for a 3.1 ms median jitter). It is this repetitive “first spike wave” that STDP is able to catch. Why does this happen? Each time the postsynaptic neuron fires to the repeating wave (by chance at the start of the learning phase), STDP reinforces the synaptic connections with the afferents that took part in firing it, thereby increasing the probability of firing again next time the wave comes. This “reinforcement of causality link” causes the selectivity to emerge (Masquelier et al., 2008), as we shall see below.

Figure 3*B* plots the membrane potential of the downstream neuron. At this stage it is not yet selective to the pattern (in gray), which is not surprising since we started from uniformly distributed random weights (Fig. 3*C*). For the same reason, the pattern activation levels and the corresponding weights are uncorrelated at this stage (Fig. 3*D*).

Figure 4 plots the same variables after 802.5 s of simulated time and demonstrates that selectivity has emerged. Postsynaptic spikes tend to indicate periods when the pattern is present, but detection only occurs when there is a reset during the pattern. The final weight distribution is bimodal, as can be seen in Figure 4*C*, with synapses either fully depressed or maximally reinforced, as usual with additive STDP (van Rossum et al., 2000). At this stage the system has converged and the synaptic weights are stable. As can be seen from Figure 4*D*, the synapses that are reinforced correspond to afferents within the pattern with the highest activation levels, which in turn correspond to the first spikes in the wave illustrated in the insets of Figure 3*A*. This is not surprising, since when a neuron is presented successively with similar waves of input spikes STDP is known to have the effect of concentrating synaptic weights on afferents that consistently fire early (Song et al., 2000; Delorme et al., 2001; Gerstner and Kistler, 2002; Guyonneau et al., 2005; Guyonneau, 2006; Masquelier and Thorpe, 2007).

Using a STDP window with a negative integral like here [i.e., $\tau^+ a^+ < \tau^- a^-$, that is, LTD tends to overcome long-term potentiation (LTP)] leads to a subthreshold regime where the postsynaptic neuron performs coincidence detection (Kempler et al., 2001). Here, it became sensitive to the nearly simultaneous arrival of the earliest spikes of the wave of Figure 3*A* insets by reinforcing the corresponding synapses. When the pattern is presented, these earliest coincident spikes tend to make the postsynaptic neuron fire. The postsynaptic spikes thus tend to indicate the presence of the pattern, and the robustness of the detection can be quantified with mutual information (see Materials and Methods).

We thus continued the simulation until $t = 1000$ s and computed the mutual information over the [800 s, 1000 s] period. This was done for 10 identical simulations, with different pseudorandomly generated input matrices and initial weights. The average mutual information was ~ 0.3 bits. We then varied x , and computed the

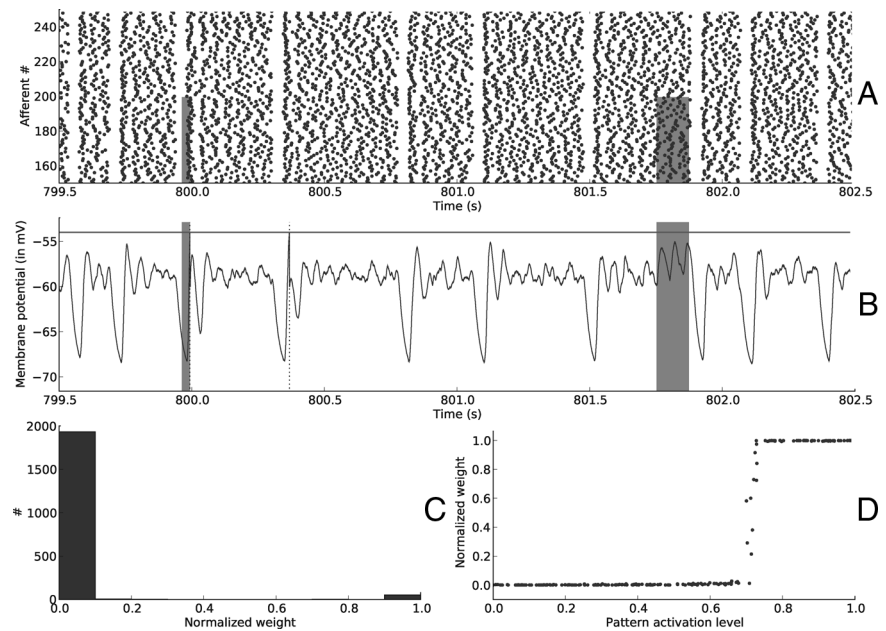


Figure 4. Resets—end of learning. Here we plotted the same variables as in Figure 3 after ~ 800 s of learning. *A*, Input spike trains. *B*, Postsynaptic membrane potential as a function of time. Selectivity has emerged. Resets still provoke oscillations in the membrane potential, but they are usually too weak to reach the threshold, unless the reset was performed during a pattern period (as it is the case at $t \approx 800$ s). This is not always the case, since at $t \approx 800.4$ s there is a false alarm. The second pattern presentation is missed because there was no reset during the pattern and therefore the spike times do not match those of the repeating spike waves of Figure 3*A* insets. *C*, Distribution of synaptic weights at $t = 802.5$ s. It became bimodal, with ~ 60 synapses maximally reinforced and the rest of them fully depressed. *D*, Synaptic weights for afferents 0...199 (involved in the pattern) as a function of their corresponding pattern activation levels. STDP reinforced the afferents with the highest levels.

average mutual information each time (note that the values for I_{\max} and a^-/a^+ were reoptimized for each x value, using an exhaustive search procedure, as described above). The results are shown in Figure 7 (dashed line). The curve has an asymptote at ~ 0.4 bits. At this stage increasing the proportion of afferents involved in the pattern does not help much. What limits the detector’s performance is mainly that the patterns are missed if no reset occurs during the presentations. Thus better performance can be achieved by using more frequent resets, as we verified by running another batch of simulations with resets every 125 ms on average (dotted line). However, we were mainly interested in the shape of the curves, and it is surprising that the performance cutoff only occurs at $x \approx 10\%$.

Benchmark 4: LIF input neurons + oscillatory drive

We then suppressed the resets, and plugged an oscillatory current (Fig. 2, Eq. 2) and affinely mapped the matrix activation levels into an additional static input current between $0.95 \cdot I_{\text{thr}}$ and $1.07 \cdot I_{\text{thr}}$. This leads to a current-to-phase conversion: the neurons that receive the strongest static currents will fire first during the phase of the cycle (Buzsáki and Chrobak, 1995; Hopfield, 1995; Mehta et al., 2002; Brody and Hopfield, 2003; Buzsáki and Draguhn, 2004; Lisman, 2005; Fries et al., 2007). Here the static current range is such that each afferent emits between one and three spikes per cycle, leading to a mean input spike rate of 14.2 Hz. The phase of the first spike is a decreasing function of the static input current. What is remarkable is the speed of convergence: with constant static currents phase locking occurs in 1–2 cycles, regardless of the initial membrane potential. The phase thus rapidly depends almost entirely on the current input current, and not on the neuron’s past history.

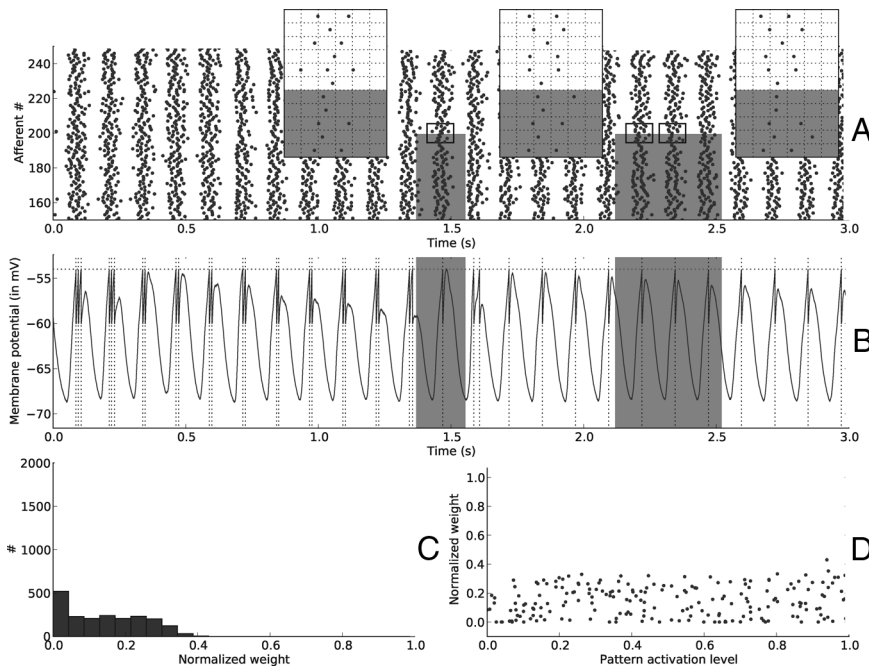


Figure 5. Oscillations—beginning of learning. Here we plotted the same variables as in Figure 3, but the afferents now receive an oscillatory drive. **A**, Input spike trains. Because of the oscillatory drive input spikes come in waves. Again, pattern periods are shown in gray. Three insets zoom on adequate periods to illustrate that the spike phases of the afferents involved in the pattern (0 ... 199) are the same (except for the noise) for different pattern presentations [the phase grid has a 1 rad (20 ms) step]. This is not true for afferents not involved in the pattern (200 ... 1999) (however those afferents have the same phases in the two right insets because their activation levels did not change). It is this repeating “spike wave” that STDP can detect and learn. **B**, Postsynaptic membrane potential as a function of time. It oscillates, since input spikes come in waves. At this time the postsynaptic neuron is not selective: the threshold is reached at least once at each cycle, whether or not the pattern is present. **C**, Distribution of synaptic weights at $t = 3$ s. Weights are still almost uniformly distributed at this stage, since few STDP updates have been made. **D**, Synaptic weights for afferents 0 ... 199 (involved in the pattern) as a function of their corresponding pattern activation levels. Both variables are uncorrelated at this stage.

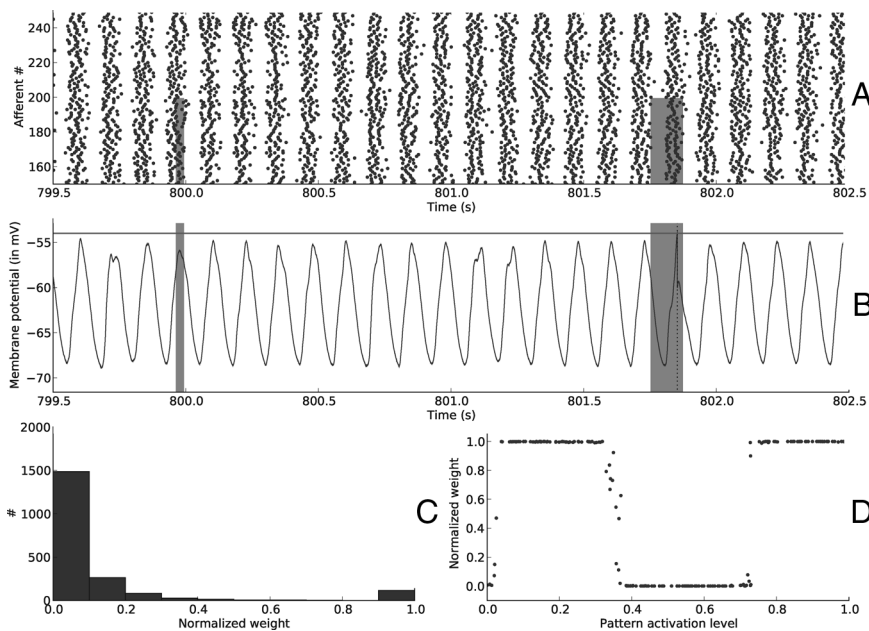


Figure 6. Oscillations—end of learning. Here we plotted the same variables as in Figure 5 after ~ 800 s of learning. **A**, Input spike trains. **B**, Postsynaptic membrane potential as a function of time. It still oscillates, but now selectivity has emerged: the threshold is reached if and only if the pattern is present for a sufficient time, in the order of one oscillation period (here 125 ms). The first pattern presentation, which lasts only 30 ms, is missed. **C**, Distribution of synaptic weight at $t = 802.5$ s. It became bimodal, with ~ 130 synapses maximally reinforced and the rest of them fully depressed. **D**, Synaptic weights for afferents 0 ... 199 (involved in the pattern) as a function of their corresponding pattern activation levels. The function is piece-wise constant (see text).

We again set $x = 10\%$, and $I_{\max} = 0.05$ nA (meaning the equivalent of 35 synchronous spikes arriving through maximally reinforced synapses are needed to reach the threshold from the resting state) and $a^-/a^+ = 1.48$, after the exhaustive 2D grid search described above.

Figure 5 describes the situation at the beginning of the simulation, and should be compared with Figure 3. Figure 5A plots the input spike trains. It can be seen that most of the spikes are emitted during the rising part of the sinusoid. Three insets zoom on adequate periods to demonstrate that the oscillations within patterns generate identical spike waves (except for the noise, here responsible for a 1.1 ms median jitter) for the afferents involved in the pattern. It is this repeating spike wave that STDP is going to catch, again thanks to its “reinforcement of causality links” property.

Figure 5B plots the postsynaptic membrane potential. At this stage the postsynaptic neuron is not yet selective: it fires at least once at every cycle, whether or not the pattern is there. This is not surprising since we start from uniformly distributed random weights (Fig. 5C). For the same reason, the pattern activation levels and the corresponding weights are uncorrelated at this stage (Fig. 5D).

Again, Figure 6 plots the same variables after 802.5 s of simulated time and demonstrates that selectivity has emerged: postsynaptic spikes are now related to the presence of the pattern. The weights have converged, again toward a bimodal distribution (Fig. 6C), with ~ 130 selected synapses out of 2000.

Figure 6D plots the relationship between the pattern activation levels and the final synaptic weights for the 200 afferents involved in the pattern. In this simulation after convergence the phase of the postsynaptic spike when the pattern is presented is ~ 5.1 rad. Very weak pattern activation levels (i.e., input currents) lead to one presynaptic spike per cycle with a phase > 5.1 rad. The corresponding synapses are thus systematically depressed, until zero is reached, as can be seen in the left part of the graph. Slightly greater activation levels lead to one presynaptic spike per cycle with phase lower than 5.1 rad; the corresponding synapses are thus fully reinforced. Even greater levels lead so several spikes per cycle with a combination of LTP and LTD, and one of them overcomes the other depending on the detailed relationships between presynaptic and postsynaptic spike phases. Note that other static input current ranges may lead




Figure 7. Mutual information between the postsynaptic spikes and the presence of the pattern, for the reset (at 4 and 8 Hz on average) and the oscillation cases (at 8 Hz) as a function of the proportion of afferents in the pattern (x). For all mechanisms the performance cutoff only occurs at $x \approx 10\%$.

to at most 1 spike per cycle—in this simpler case STDP ends up reinforcing the synapses with the highest pattern activation levels, as in the reset case.

Again, the fact that we used an STDP window with a negative integral led to a subthreshold regime where the postsynaptic neuron performs coincidence detection (Kempster et al., 2001). Here, it became sensitive to the nearly simultaneous arrival of a subset of the spikes within the wave of Figure 5A insets and again postsynaptic spikes tend to indicate the presence of the pattern.

As before, we quantified the detection performance by continuing the simulation until $t = 1000$ s and computing the mutual information over the [800 s, 1000 s] period, using exactly the same procedure. This was done for various x values, running each time 10 identical simulations, with different pseudorandomly generated input matrices and initial weights. The values for I_{\max} and a^-/a^+ were reoptimized for each x value. The results are shown in Figure 7 (solid line). The curve has an asymptote at ~ 0.55 bits. At this stage increasing the number of involved afferents again does not help much. What limits the detector performance is mainly that the patterns are missed if they are shown for less than one period. Performance is slightly higher than with resets at the same frequency. However, we were again more interested in the shape of the curve, and it is surprising that the performance cutoff only occurs at $x \approx 10\%$.

For what frequency range does the proposed mechanism work? Of course the problem is invariant to a scaling of all the time constants involved. These are of four kinds: the STDP time constants (τ^+ and τ^-), the neuronal and synaptic time constants (τ_m , τ_s , and refractory period), the time constants that are specific to the problem (mean Δt and mean inter pattern interval), and the oscillation period. Scaling the STDP time constants does not make much sense, since most of the experimentalists find $\tau^+ \approx 20$ –30 ms and $\tau^- \approx 10$ –20 ms (Caporale and Dan, 2008). Scaling the neuronal time constants makes more sense, since a broad range of values are found in the brain, and the neurons can be in a high-conductance state (Destexhe et al., 2003). Scaling the problem time constant is also legitimate, since the brain has to deal with signals with a broad range of intrinsic timescales. We




Figure 8. Acceptable frequency range. Here we plot the mutual information as a function of the oscillation frequency, for two STDP modes: the “all-to-all” spike interaction mode, used in the baseline simulations, has a performance cutoff at ~ 25 Hz. Another commonly used mode, the “nearest spikes” mode extends the cutoff to ~ 40 Hz. Note that the two modes converge for slow oscillations, because spikes are so distant that only the nearest will have a significant impact.

thus kept the STDP time window constant and varied the oscillation period, rescaling the neuronal and problem time constants accordingly (note that this is numerically equivalent to scaling only the STDP window while keeping everything else constant). This means in particular that the inputs for the STDP neuron were the same except for a time scaling factor. By scaling the neuron time constants with the oscillation period, we avoid taking into account signal detection aspects (for example, a neuron with a high τ_m could not discriminate between two different short spike waves). By scaling the problem time constants with the oscillation period, we counterbalance the advantage of fast oscillations for detecting short pattern presentations. This way we only study the interactions between the oscillations and the STDP intrinsic time scales, all other things being equal.

Figure 8 plots the results in terms of mutual information estimated over the [800 s, 1000 s] period for 10 simulations with $x = 10\%$. The values for I_{\max} and a^-/a^+ were optimized for each point using the exhaustive search procedure described earlier, though this time the ratio for the a^-/a^+ was $1.05^{0.5} \approx 1.02$, as it needed to be fine-tuned with fast oscillations. The “all-to-all” mode, used in the baseline simulations, gave acceptable results up to ~ 25 Hz. But with faster oscillations, the weight updates are too corrupted by preceding and subsequent spike waves and learning fails. A way to limit this problem is to implement the “nearest spikes” mode. This shifts the curve of 0.5 to 1 octave, leading to a cutoff of ~ 40 Hz. However, there is still much debate about whether “nearest spikes” or “all-to-all” modification rules for STDP are more biologically realistic (Burkitt et al., 2004).

Discussion

Several authors have already proposed oscillation-based mechanisms that can perform a current-to-phase conversion (Buzsáki and Chrobak, 1995; Hopfield, 1995; Mehta et al., 2002; Brody and Hopfield, 2003; Buzsáki and Draguhn, 2004; Lisman, 2005; Fries et al., 2007). Recently, such mechanisms have received experimental support *in vitro* (McLelland and Paulsen, 2009), and could account for the numerous cases of PoFC observed *in vivo* (see Table 1). More generally, there is both theoretical (Brette and Guigon, 2003) and experimental (Hasenstaub et al., 2005; Schaefer et al., 2006; Markowitz et al., 2008) evidence that a common oscillatory drive for a group of neurons, periodic or not, improves the reliability of their spike times, by decreasing their

sensitivity to the initial conditions, and avoiding jitter accumulation, so that the spike times depend reliably on the current input values.

But this raises the question of why the brain would need to use such an oscillatory mechanism to control spike times. We suggest that it is to facilitate downstream processing by STDP. Surprisingly, a single neuron equipped with the learning rule robustly detects and learns a repeating activation pattern that affects an unknown subset of only 10% of its afferents, under conditions where the duration of the patterns is unpredictable and where the patterns recur at unpredictable intervals. After learning, recognition is robust, and only takes one cycle. More conventional rate-based coding/learning schemes would require at least a few interspike intervals to estimate the rates (Gautrais and Thorpe, 1998), and a rate-based Hebbian learning rule instead of STDP. This is the reason why benchmarks 1 and 2, in which only the rates repeat, failed.

The use of global potential resets is a similar alternative to a common oscillatory drive, which also leads to reproducible spike times that STDP can pick. This may be a valid description of what happens during discrete sensory processing such as saccades or sniffs (Uchida et al., 2006), or with “stimulus onset paradigms.” But the fact that periods of oscillatory activity are found throughout the brain, and the suggestion that they could be particularly useful for the continuous cognitive processes, means that the oscillatory drive model may be of considerable theoretical importance.

Consistent with this model, a growing body of experimental evidence in animals and humans demonstrates that successful long-term memory encoding correlates with increased oscillatory activity across a broad range of frequencies (from theta to gamma), in both sensory and associative areas [for recent reviews, see Jensen et al. (2007), Klimesch et al. (2008), and Tallon-Baudry (2009)]. Interestingly, beyond mere oscillation power what seems to be a prerequisite for successful memory formation is that single units should be phase-locked to the oscillation (Rutishauser et al., 2009), as in our model. In addition, even prestimulus oscillatory activity can predict successful episodic memory encoding in humans (Guderian et al., 2009), indicating that the facilitating oscillations are ongoing and internally generated, and *not* evoked by the stimulus—as the ones we used in this work.

How about PoFC in the high gamma range (say >40 Hz)? Our results suggest that it is not optimal for STDP-based learning and decoding, at least with the STDP time windows observed so far, but not that it does not exist. However evidence for it is still scarce in the cortex. For example, Montemurro et al. (2008), Ray et al. (2008), and Kayser et al. (2009) all looked for it unsuccessfully [but see König et al. (1995), Fries et al. (2001), Hoffman et al. (2009), and Koepsell et al. (2009) for exceptions]. In the studies by Jacobs et al. (2007) and Ray et al. (2008), some spikes did lock to gamma oscillations, but the preferred phases always tended to be near the oscillation peaks, ruling out PoFC, but suggesting instead a binary code in which information would be coded in the combination of neurons which fire at least once in each oscillation cycle. Such a code could probably be decoded by STDP, but it is beyond the scope of this paper.

The ability of STDP to detect and learn repeating spike patterns had been noted before in the case of continuous activity (Masquelier et al., 2008). Oscillations are thus not needed for the proposed mechanism to work provided some input spike patterns repeat reliably. But first, as mentioned earlier, oscillations provide one appealing way to generate the repeating precise input

spike times [an issue that was not addressed by Masquelier et al. (2008)]. Second, spike patterns appear to be easier to learn when embedded in oscillatory activity than in continuous background activity: learning becomes possible when only 10% of the afferents are involved in the repeating patterns, compared with a minimum value of about 40% for continuous activity in the previous study (Masquelier et al., 2008). Note, however, that in both studies the values refer to the minimum number of afferents that need to be involved to allow detection by a single neuron. Reliability can easily be increased by increasing the number of “listening” neurons. The mechanism using an oscillatory drive is also more robust to changes of parameters, in particular of the threshold.

Somewhat surprisingly, only a few papers have studied STDP in the context of oscillating presynaptic activity. One of them (Gerstner et al., 1996) showed how STDP selects connections with matching delays from a broad distribution of axons with random delays, giving rise to fast and time-accurate postsynaptic responses. Yoshioka (2002) also demonstrated the pertinence of STDP in oscillating associative memory networks. Closer to our work, in a recent study modeling the locust olfactory system, the authors also injected the output of an oscillating neural population into a neuron equipped with STDP, and convincingly demonstrated that the rule did a better job than a more conventional rate-based Hebbian rule at explaining both the sparsity and the selectivity observed experimentally (Finelli et al., 2008). They used a discrete finite set of (odor) stimuli, and each of them had to be presented several times in a row for the system to learn it. This is quite different from what we are doing here, namely unsupervised learning of an arbitrary pattern of real valued activation levels, that repeatedly appears during unpredictable durations and at unpredictable intervals, and that does not concern all the afferents.

What happens if there is more than one repeating pattern present in the input? We verified that as the learning progresses, the increasing selectivity of the postsynaptic neuron rapidly prevents it from responding to several patterns. Instead, it randomly picks one, and becomes selective to it and only to it. To learn the other patterns other neurons listening to the same inputs are needed. As in our previous work (Masquelier et al., 2009), lateral inhibitory connections between those “listening” neurons could implement a competitive mechanism in which the first neuron to fire strongly inhibits the others, thereby preventing the neurons from learning the same patterns.

Together these results suggest how two simple mechanisms present in the brain may combine to induce a kind of temporal coding: oscillations allow information to be encoded in the spike phases, and STDP provides an appealing mechanism that can learn how to decode it. The two mechanisms interact constructively for a large range of oscillation frequencies. Of course, we do not claim this code is the only one at work in the brain. Population rate coding, temporal coding with respect to stimulus onset and PoFC, to cite only these schemes, are not mutually exclusive, but could be nested to encode information on different time scales (Kayser et al., 2009).

References

- Albrecht DG, Geisler WS, Frazor RA, Crane AM (2002) Visual cortex neurons of monkeys and cats: temporal dynamics of the contrast response function. *J Neurophysiol* 88:888–913.
- Amarasingham A, Chen TL, Geman S, Harrison MT, Sheinberg DL (2006) Spike count reliability and the Poisson hypothesis. *J Neurosci* 26:801–809.
- Bi G, Poo M (2001) Synaptic modification by correlated activity: Hebb's postulate revisited. *Annu Rev Neurosci* 24:139–166.

- Brette R, Guigon E (2003) Reliability of spike timing is a general property of spiking model neurons. *Neural Comput* 15:279–308.
- Brody CD, Hopfield JJ (2003) Simple networks for spike-timing-based computation, with application to olfactory processing. *Neuron* 37:843–852.
- Burkitt AN, Meffin H, Grayden DB (2004) Spike-timing-dependent plasticity: the relationship to rate-based learning for models with weight dynamics determined by a stable fixed point. *Neural Comput* 16:885–940.
- Buzsáki G, Chrobak JJ (1995) Temporal structure in spatially organized neuronal ensembles: a role for interneuronal networks. *Curr Opin Neurobiol* 5:504–510.
- Buzsáki G, Draguhn A (2004) Neuronal oscillations in cortical networks. *Science* 304:1926–1929.
- Caporale N, Dan Y (2008) Spike timing-dependent plasticity: a hebbian learning rule. *Annu Rev Neurosci* 31:25–46.
- Cassenaer S, Laurent G (2007) Hebbian STDP in mushroom bodies facilitates the synchronous flow of olfactory information in locusts. *Nature* 448:709–713.
- Celebrini S, Thorpe S, Trotter Y, Imbert M (1993) Dynamics of orientation coding in area V1 of the awake primate. *Vis Neurosci* 10:811–825.
- Chase SM, Young ED (2007) First-spike latency information in single neurons increases when referenced to population onset. *Proc Natl Acad Sci U S A* 104:5175–5180.
- David FO, Hugues E, Cenier T, Fourcaud-Trocmé N, Buonviso N (2009) Specific entrainment of mitral cells during gamma oscillation in the rat olfactory bulb. *PLoS Comput Biol*, in press.
- Delorme A, Perrinet L, Thorpe J, Samuelides M (2001) Networks of integrate-and-fire neurons using rank order coding B: spike timing dependent plasticity and emergence of orientation selectivity. *Neurocomputing* 38–40:539–545.
- Destexhe A, Rudolph M, Paré D (2003) The high-conductance state of neocortical neurons in vivo. *Nat Rev Neurosci* 4:739–751.
- Eeckman FH, Freeman WJ (1990) Correlations between unit firing and EEG in the rat olfactory system. *Brain Res* 528:238–244.
- Finelli LA, Haney S, Bazhenov M, Stopfer M, Sejnowski TJ (2008) Synaptic learning rules and sparse coding in a model sensory system. *PLoS Comput Biol* 4:e1000062.
- Friedrich RW, Habermann CJ, Laurent G (2004) Multiplexing using synchrony in the zebrafish olfactory bulb. *Nat Neurosci* 7:862–871.
- Fries P, Neuenschwander S, Engel AK, Goebel R, Singer W (2001) Rapid feature selective neuronal synchronization through correlated latency shifting. *Nat Neurosci* 4:194–200.
- Fries P, Nikolaić D, Singer W (2007) The gamma cycle. *Trends Neurosci* 30:309–316.
- Gautrais J, Thorpe S (1998) Rate coding versus temporal order coding: a theoretical approach. *Biosystems* 48:57–65.
- Gawne TJ, Kjaer TW, Richmond BJ (1996) Latency: another potential code for feature binding in striate cortex. *J Neurophysiol* 76:1356–1360.
- Gerstner W, Kistler W (2002) *Spiking neuron models*. Cambridge, UK: Cambridge UP.
- Gerstner W, Kempter R, van Hemmen JL, Wagner H (1996) A neuronal learning rule for sub-millisecond temporal coding. *Nature* 383:76–81.
- Gollisch T, Meister M (2008) Rapid neural coding in the retina with relative spike latencies. *Science* 319:1108–1111.
- Goodman D, Brette R (2008) Brian: a simulator for spiking neural networks in python. *Front Neuroinformatics* 2:5.
- Guderian S, Schott BH, Richardson-Klavehn A, Düzel E (2009) Medial temporal theta state before an event predicts episodic encoding success in humans. *Proc Natl Acad Sci U S A* 106:5365–5370.
- Gütig R, Aharonov R, Rotter S, Sompolinsky H (2003) Learning input correlations through nonlinear temporally asymmetric hebbian plasticity. *J Neurosci* 23:3697–3714.
- Guyonneau R (2006) Codage par latence et STDP: des stratégies temporelles pour expliquer le traitement visuel rapide. PhD thesis, Université Toulouse III-Paul Sabatier.
- Guyonneau R, VanRullen R, Thorpe SJ (2005) Neurons tune to the earliest spikes through STDP. *Neural Comput* 17:859–879.
- Hafting T, Fyhn M, Bonnevie T, Moser MB, Moser EI (2008) Hippocampus-independent phase precession in entorhinal grid cells. *Nature* 453:1248–1252.
- Hasenstaub A, Shu Y, Haider B, Kraushaar U, Duque A, McCormick DA (2005) Inhibitory postsynaptic potentials carry synchronized frequency information in active cortical networks. *Neuron* 47:423–435.
- Hoffman K, Turesson H, Logothetis N (2009) Phase coding of faces and objects in the superior temporal sulcus. *Frontiers in Computational Neuroscience Conference Abstract: Computational and Systems Neuroscience*.
- Hopfield JJ (1995) Pattern recognition computation using action potential timing for stimulus representation. *Nature* 376:33–36.
- Jacobs AL, Fridman G, Douglas RM, Alam NM, Latham PE, Prusky GT, Nirenberg S (2009) Ruling out and ruling in neural codes. *Proc Natl Acad Sci U S A* 106:5936–5941.
- Jacobs J, Kahana MJ, Ekstrom AD, Fried I (2007) Brain oscillations control timing of single-neuron activity in humans. *J Neurosci* 27:3839–3844.
- Jensen O, Kaiser J, Lachaux JP (2007) Human gamma-frequency oscillations associated with attention and memory. *Trends Neurosci* 30:317–324.
- Johansson RS, Birznieks I (2004) First spikes in ensembles of human tactile afferents code complex spatial fingertip events. *Nat Neurosci* 7:170–177.
- Kashiwadani H, Sasaki YF, Uchida N, Mori K (1999) Synchronized oscillatory discharges of mitral/tufted cells with different molecular receptive ranges in the rabbit olfactory bulb. *J Neurophysiol* 82:1786–1792.
- Kayser C, Montemurro MA, Logothetis NK, Panzeri S (2009) Spike-phase coding boosts and stabilizes information carried by spatial and temporal spike patterns. *Neuron* 61:597–608.
- Kempter R, Gerstner W, van Hemmen JL (1999) Hebbian learning and spiking neurons. *Phys Rev E* 59:4498–4514.
- Kempter R, Gerstner W, van Hemmen JL (2001) Intrinsic stabilization of output rates by spike-based hebbian learning. *Neural Comput* 13:2709–2741.
- Kiani R, Esteky H, Tanaka K (2005) Differences in onset latency of macaque inferotemporal neural responses to primate and non-primate faces. *J Neurophysiol* 94:1587–1596.
- Klimesch W, Freunberger R, Sauseng P, Gruber W (2008) A short review of slow phase synchronization and memory: evidence for control processes in different memory systems? *Brain Res* 1235:31–44.
- Koepsell K, Wang X, Vaingankar V, Wei Y, Wang Q, Rathbun DL, Usrey WM, Hirsch JA, Sommer FT (2009) Retinal oscillations carry visual information to cortex. *Front Syst Neurosci* 3:4.
- König P, Engel AK, Roelfsema PR, Singer W (1995) How precise is neuronal synchronization? *Neural Computation* 7:469–485.
- Lee H, Simpson GV, Logothetis NK, Rainer G (2005) Phase locking of single neuron activity to theta oscillations during working memory in monkey extrastriate visual cortex. *Neuron* 45:147–156.
- Lisman J (2005) The theta/gamma discrete phase code occurring during the hippocampal phase precession may be a more general brain coding scheme. *Hippocampus* 15:913–922.
- Liu RC, Tzonev S, Rebrink S, Miller KD (2001) Variability and information in a neural code of the cat lateral geniculate nucleus. *J Neurophysiol* 86:2789–2806.
- Markowitz DA, Collman F, Brody CD, Hopfield JJ, Tank DW (2008) Rate-specific synchrony: using noisy oscillations to detect equally active neurons. *Proc Natl Acad Sci U S A* 105:8422–8427.
- Masquelier T, Thorpe SJ (2007) Unsupervised learning of visual features through spike timing dependent plasticity. *PLoS Comput Biol* 3:e31.
- Masquelier T, Guyonneau R, Thorpe SJ (2008) Spike timing dependent plasticity finds the start of repeating patterns in continuous spike trains. *PLoS ONE* 3:e1377.
- Masquelier T, Guyonneau R, Thorpe SJ (2009) Competitive STDP-based spike pattern learning. *Neural Comput* 21:1259–1276.
- McLelland D, Paulsen O (2009) Neuronal oscillations and the rate-to-phase transform: mechanism, model and mutual information. *J Physiol* 587:769–785.
- Mehta MR, Lee AK, Wilson MA (2002) Role of experience and oscillations in transforming a rate code into a temporal code. *Nature* 417:741–746.
- Montemurro MA, Rasch MJ, Murayama Y, Logothetis NK, Panzeri S (2008) Phase-of-firing coding of natural visual stimuli in primary visual cortex. *Curr Biol* 18:375–380.
- O'Keefe J, Recce ML (1993) Phase relationship between hippocampal place units and the eeg theta rhythm. *Hippocampus* 3:317–330.
- Perez-Orive J, Mazor O, Turner GC, Cassenaer S, Wilson RI, Laurent G

- (2002) Oscillations and sparsening of odor representations in the mushroom body. *Science* 297:359–365.
- Pouget A, Dayan P, Zemel R (2000) Information processing with population codes. *Nat Rev Neurosci* 1:125–132.
- Ray S, Hsiao SS, Crone NE, Franaszczuk PJ, Niebur E (2008) Effect of stimulus intensity on the spike-local field potential relationship in the secondary somatosensory cortex. *J Neurosci* 28:7334–7343.
- Rutishauser U, Mamelak A, Ross I, Schuman E (2009) Predictors of successful memory encoding in the human hippocampus and amygdala. *Frontiers in Computational Neuroscience Conference Abstract: Computational and Systems Neuroscience*.
- Schaefer AT, Angelo K, Spors H, Margrie TW (2006) Neuronal oscillations enhance stimulus discrimination by ensuring action potential precision. *PLoS Biol* 4:e163.
- Shadlen MN, Newsome WT (1998) The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci* 18:3870–3896.
- Softky WR, Koch C (1993) The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J Neurosci* 13:334–350.
- Song S, Miller KD, Abbott LF (2000) Competitive hebbian learning through spike-timing-dependent synaptic plasticity. *Nat Neurosci* 3:919–926.
- Stopfer M, Bhagavan S, Smith BH, Laurent G (1997) Impaired odour discrimination on desynchronization of odour-encoding neural assemblies. *Nature* 390:70–74.
- Tallon-Baudry C (2009) The roles of gamma-band oscillatory synchrony in human visual cognition. *Front Biosci* 14:321–332.
- Thorpe S, Imbert M (1989) Biological constraints on connectionist modeling. In: *Connectionism in perspective* (Pfeifer R, Fogelman-Soulie F, Steels L, Schreter Z, eds), pp. 63–92. Amsterdam: Elsevier.
- Uchida N, Kepecs A, Mainen ZF (2006) Seeing at a glance, smelling in a whiff: rapid forms of perceptual decision making. *Nat Rev Neurosci* 7:485–491.
- Uzzell VJ, Chichilnisky EJ (2004) Precision of spike trains in primate retinal ganglion cells. *J Neurophysiol* 92:780–789.
- van Rossum MC, Bi GQ, Turrigiano GG (2000) Stable hebbian learning from spike timing-dependent plasticity. *J Neurosci* 20:8812–8821.
- Yoshioka M (2002) Spike-timing-dependent learning rule to encode spatio-temporal patterns in a network of spiking neurons. *Phys Rev E Stat Nonlin Soft Matter Phys* 65:011903.