

Competitive STDP-Based Spike Pattern Learning

Timothée Masquelier

timothee.masquelier@alum.mit.edu

Centre de Recherche Cerveau et Cognition, Université Toulouse 3, Centre National de la Recherche Scientifique, Faculté de Médecine de Rangueil, Toulouse 31062, France, and Computational Neuroscience Group, Department of Information and Communication Technologies, Universitat Pompeu Fabra, Barcelona E-08003, Spain

Rudy Guyonneau

rg272@georgetown.edu

Centre de Recherche Cerveau et Cognition, Université Toulouse 3, Centre National de la Recherche Scientifique, Faculté de Médecine de Rangueil, Toulouse 31062, France

Simon J. Thorpe

Simon.Thorpe@cerco.ups-tlse.fr

Centre de Recherche Cerveau et Cognition, Université Toulouse 3, Centre National de la Recherche Scientifique, Faculté de Médecine de Rangueil, Toulouse 31062, France, and SpikeNet Technology SARL, Prologue 1 La Pyrénéenne, Labège 31673, France

Recently it has been shown that a repeating arbitrary spatiotemporal spike pattern hidden in equally dense distracter spike trains can be robustly detected and learned by a single neuron equipped with spike-timing-dependent plasticity (STDP) (Masquelier, Guyonneau, & Thorpe, 2008). To be precise, the neuron becomes selective to successive coincidences of the pattern.

Here we extend this scheme to a more realistic scenario with multiple repeating patterns and multiple STDP neurons “listening” to the incoming spike trains. These “listening” neurons are in competition: as soon as one fires, it strongly inhibits the others through lateral connections (one-winner-take-all mechanism). This tends to prevent the neurons from learning the same (parts of the) repeating patterns, as shown in simulations. Instead, the population self-organizes, trying to cover the different patterns or coding one pattern by the successive firings of several neurons, and a powerful distributed coding scheme emerges.

Taken together, these results illustrate how the brain could easily encode and decode information in the spike times, a theory referred to as temporal coding, and how STDP could play a key role by detecting repeating patterns and generating selective response to them.

1 Introduction

Spike-timing-dependent plasticity (STDP) is now a reasonably well-established physiological mechanism of activity-driven synaptic regulation. It has been observed extensively *in vitro* for more than a decade (Markram, Lübke, Frotscher, & Sakmann, 1997; Bi & Poo, 1998; Zhang, Tao, Holt, Harris, & Poo, 1998; Feldman, 2000) and more recently *in vivo* in the *Xenopus* visual system (Vislay-Meltzer, Kampff, & Engert, 2006; Mu & Poo, 2006), the locust mushroom body (Cassenaer & Laurent, 2007), and rat visual (Meliza & Dan, 2006) and barrel (Jacob, Brasier, Erchova, Feldman, & Shulz, 2007) cortex. It has also been shown to do a better job than more conventional Hebbian correlation-based plasticity at explaining both cortical reorganization in cat primary visual cortex (Young et al., 2007) and connectivity in locust olfactory system (Finelli, Haney, Bazhenov, Stopfer, & Sejnowski, 2008).

We recently demonstrated that a single leaky integrate-and-fire (LIF) neuron equipped with STDP is able to solve a difficult computational problem: to robustly detect a repeating arbitrary spatiotemporal spike pattern among its afferents by becoming sensitive to the successive coincidences in the pattern and tracking back through the pattern until the beginning is found. This is true even when the pattern is embedded in equally dense distracter spike trains and even when it does not involve all the afferents (Masquelier, Guyonneau, & Thorpe, 2008). The phenomenon emerges from STDP: the rule specifically reinforces connections from afferents that helped the neuron reach its threshold, so each time the neuron fires to the pattern, it increases the probability that the threshold is reached again and earlier next time the pattern is presented. As a result, the neuron becomes selective to the pattern and tracks back through it until the beginning is found. The mechanism was shown to be robust to a wide range of perturbations: missing spikes, jitter, noisy initial weights, infrequent patterns, and reductions in the proportion of afferents involved in the pattern. While this result is of theoretical importance, it is clear that a more realistic situation would be to have multiple repeating patterns present in the input and multiple neurons listening to them. This is the kind of situation we explore in this letter.

We simulated a population of 2000 afferents firing continuously for a few hundred seconds according to a Poisson process with variable instantaneous rate that was independently generated for each afferent (see the appendix for details). We inserted into these stochastic spike trains several distinct 50 ms patterns, which repeated at random intervals, as can be seen in Figure 1 (for clarity, the figure shows only 100 afferents). The repeating patterns had roughly the same spike density as the distracter parts so as to make them invisible in terms of firing rates. The firing rate averaged over the population and estimated over 10 ms time bins has a mean of 64 Hz and a standard deviation of less than 2 Hz (this firing rate is even more constant than in the 100 afferent case of Figure 1 because of the law of large numbers).




Figure 1: Spatiotemporal spike patterns. Here we show three different 50 ms long patterns (black circles, black squares, and black diamonds) that repeat at random intervals within otherwise stochastic Poissonian activity (gray circles). Because of space constraints, we show only the first 100 afferents (out of 2000). Notice that each of the three patterns involves only half of the afferents, and the involved sets overlap. The bottom panel plots the population-averaged firing rates over 10 ms time bins (we chose 10 ms because it is the membrane time constant of the neuron used in the simulations) and demonstrates that nothing about the population firing rates can be used to characterize the periods when the pattern is present. Detecting the patterns thus requires taking the spike times into account.

Each repeating pattern involved only half of the afferents (independently randomly chosen—the selected sets were thus not disjoint, and individual afferents could be implicated in multiple patterns). We increased the difficulty still further by adding a permanent 10 Hz Poissonian spontaneous activity to all the neurons and a 1 ms jitter to the patterns.

As can be seen in Figure 2, several (three on the figure) downstream LIF neurons were connected to these 2000 afferents, with random initial synaptic weights (drawn from an equiprobable distribution over $[0,1]$). They were equipped with STDP, and each integrated the spike trains in parallel. Lateral inhibitory connections were set up between them, so that as soon as a neuron fired, it sent a strong Inhibitory postsynaptic potential (IPSP) to its neighbors. This is a biologically plausible way to implement a one-winner-take-all (1WTA) mechanism (Thorpe, 1990). We will see through

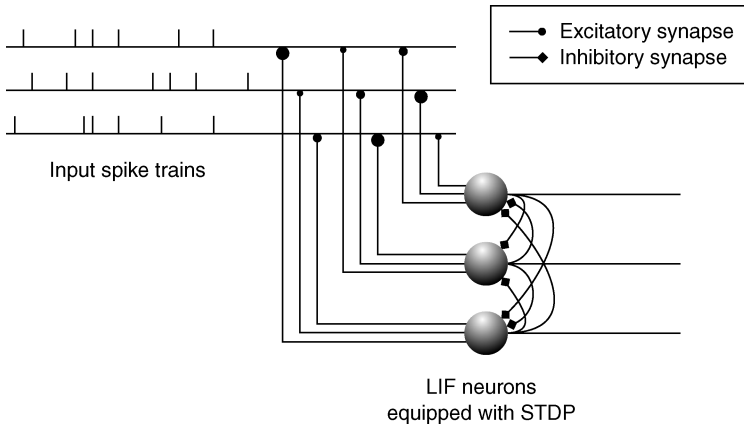


Figure 2: Network architecture. The 2000 input spike trains (we represented only 3) are integrated in parallel by several (here three) downstream LIF neurons, through excitatory synapses governed by STDP. Lateral inhibitory connections are set up between them, so that as soon as a neuron fires, it sends a strong IPSP to its neighbors (the apparent violation of Dale's law can be resolved via inhibitory interneurons). This is a biologically plausible way to implement a one-winner-take-all mechanism (Thorpe, 1990). Unlike the excitatory connections, the inhibitory connections are hard-wired and not plastic.

simulations that the neurons learn to detect the repeating patterns and that the inhibition tends to prevent them from learning the same parts of them. Instead, the neural population self-organizes, trying to cover all the repeating input patterns or to code one pattern by the successive firing of several neurons.

Note that the effect of winner-take-all on STDP-based learning has already been studied with both discrete spike volleys (Delorme, Perrinet, Thorpe, & Samuelides, 2001; Guyonneau, VanRullen, & Thorpe, 2004; Masquelier & Thorpe, 2007) and periodic inputs (Gerstner, Ritz, & van Hemmen, 1993; Yoshioka, 2002). However, this is the first time that this sort of mechanism has been used in a continuous regime in which patterns occur at random intervals.

2 Results

2.1 One Pattern, n Neurons. As a first extension to our previous study (Masquelier et al., 2008), we explored a situation in which there was still only one repeating pattern but three downstream neurons. Without inhibition, the three neurons are independent; they thus all behave exactly as in our previous study. Specifically, they all learn to detect the repeating pattern and track back through it until the beginning is found. The latencies of the

postsynaptic spikes with respect to the beginning of the pattern are then minimal (around 5 ms). Some differences between neurons may remain after convergence due to differences in the initial synaptic weights, but these variations are tiny (mean differences in postsynaptic latencies is 0.15 ms, with standard deviation 0.24 ms, estimated over 100 repetitions of a 225 second simulation sequence using three postsynaptic neurons).

The situation becomes more interesting with inhibition. As before, one neuron (A) usually finds the start of the pattern and fires to it. But if a second neuron B starts tracking back through the pattern, it cannot reach the start of the pattern: since A inhibits B rapidly after the pattern is shown, B will need some time to reach its threshold after A has fired. B thus has a postsynaptic latency that is significantly greater than A's. If a third neuron C starts tracking back through the pattern, its progression will be stopped by B. The neurons thus "stack," that is, they fire to successive parts of the pattern, and the delay between two successive parts depends on the strength of the inhibition (the amplitude of the IPSP).

To demonstrate these points, we first ran a 225 s simulation with three postsynaptic neurons and an IPSP of amplitude one-fourth of threshold (i.e. 137.5 arbitrary units). Figure 3 (top) plots the membrane potentials of the three neurons at the beginning of the simulation. Notice how the neurons inhibit each other each time they fire. At this phase of the simulation, the neurons are not selective. Since the presynaptic spike density is almost constant, the thresholds are reached pseudoperiodically. We represented the periods when the repeating pattern was presented with a gray rectangle, and it can be seen that the neurons fire inside and outside it indifferently. But as in our previous study (Masquelier et al., 2008), STDP rapidly causes the neurons to become selective to the repeating pattern. What is new here is that inhibition prevents them from all tracking back through the entire pattern until the beginning. Instead the neurons stack; they become selective to successive parts of the pattern, as can be seen in Figure 3 (middle). In Figure 3 (bottom), we represent the postsynaptic-like latencies of the three neurons with respect to the beginning of the pattern. By convention, the latency is said to be zero if the neuron fired outside the pattern, that is, generated a false alarm (these zero latencies are not taken into account in the mean latencies). Notice how the neurons track back through the pattern before stacking from the beginning of it.

We then varied the amplitude of the IPSP (0.1, 0.25, 0.5, and 1 times the threshold). For each value, we ran 100 simulations with different pseudo-randomly generated input spike trains. Most of the neurons usually learned something (73% on average), that is, became selective to one part of the pattern (criteria: a hit rate superior to 90% and a false alarm rate inferior to 1 Hz). We systematically reported the postsynaptic spike latencies for those neurons and computed the differences between two successive latencies. It can be seen in Figure 4 that as expected, the stronger the inhibition, the longer the interval between two successive firings. A given pattern

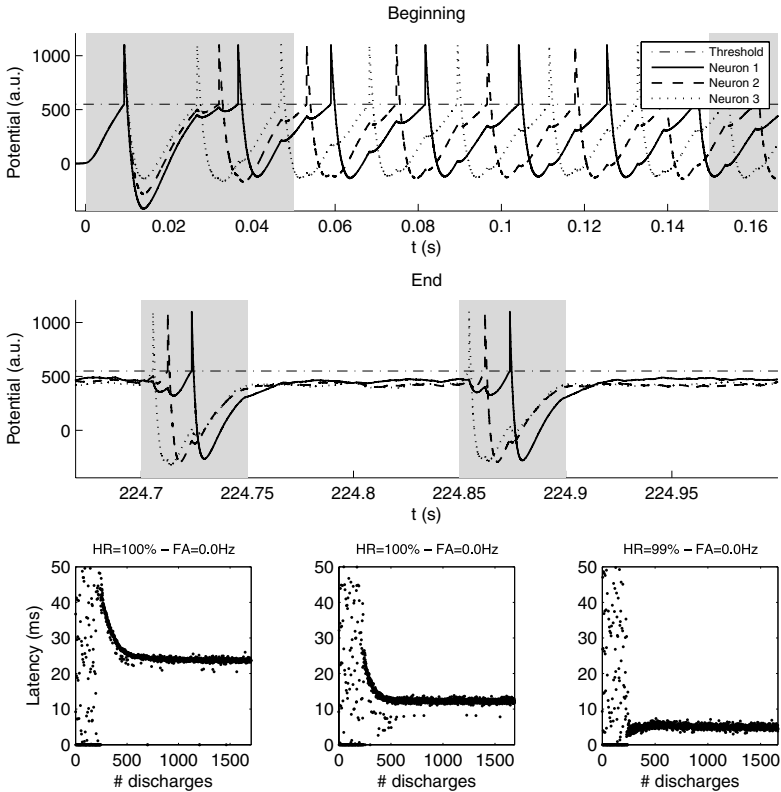


Figure 3: Learning. The top and middle panels plot the membrane potential of the three neurons (arbitrary units), as a function of time. Peaks correspond to postsynaptic spikes and are followed by a negative spike afterpotential. Notice that each time a neuron fires, it sends an IPSP to the two others. This tends to prevent neurons from firing too close to each other in time. (Top) Beginning of the simulation. The neurons are not selective. They fire pseudoperiodically, both inside and outside the pattern (shown in gray) indifferently. (Middle) End of the simulation. Each neuron has become selective to a different part of the pattern (i.e., has a different latency with respect to the beginning of the pattern; see also the bottom panel). Note that the timescale is not the same for the top and bottom plots. (Bottom) Postsynaptic spike latencies of the three neurons during the learning process. By convention, the latency is said to be zero if the neuron fired outside the pattern (i.e., generated a false alarm). Notice that these false alarms stop after about 250 discharges: selectivity has emerged. Then the neurons start tracking back through the pattern, but because of inhibition, each one is stopped in its progression by the preceding one. The neurons thus “stack” from the beginning of the pattern, with latencies of, respectively, 5, 12, and 24 ms. Hit and false alarm rates are shown above the plots (both computed over the last 75 s of simulated time).




Figure 4: Effect of inhibition on latency differences. Here we plotted the mean latency difference between two successive successful neurons (criteria: hit rate above 90%, false alarm rate below 1 Hz) for various inhibition strengths (expressed as the magnitude of the IPSP with respect to the neuron's threshold). As expected, the stronger the inhibition, the longer the intervals between two successive firings.

can thus be coded by more or fewer neurons, depending on the strength of the inhibition. To maximize coding efficiency, the system should avoid having redundant neurons coding for virtually identical temporal parts of a pattern, but should ensure the entire pattern is coded by avoiding having blind zones where the input does not influence the output. A typical delay of the order of the membrane time constant (here 10 ms) is thus close to optimal. This corresponds to an IPSP of amplitude roughly one-fourth of threshold. We used this value for the rest of the simulations.

2.2 n Patterns, m Neurons. What happens if more than one repeating pattern is present in the input? Usually the same neuron cannot become selective to two distinct patterns unless they happen to be very similar (only once in the 100 simulations described below did a neuron become responsive to two patterns). Chance determines the pattern to which each neuron becomes selective, but inhibition encourages the neurons to distribute themselves across all the patterns.

To demonstrate these points, we ran one hundred 675-second-long simulations with different pseudorandomly generated spike trains, all of them

involving three repeating patterns and nine neurons. Figure 5 shows a typical result. In this example patterns 1 and 2 were learned by two neurons, and pattern 3 by three neurons. Two neurons “died”: they stopped firing after too many discharges had occurred outside the patterns, causing too much depression (in this letter, the STDP function is biased toward depression; see the appendix).

Figure 6 shows statistics over the 100 runs. In more than two-thirds of the cases, each of the three patterns was learned by at least one neuron. On average, 5.7 out of 9 neurons (63%) learned something (criteria: hit rate over 90%, false alarm rate below 1 Hz). This may seem lower than in the single-neuron case studied previously (Masquelier et al., 2008), where the pattern was learned in 96% of the cases. However, the problem faced by a neuron B when it tries to learn a pattern that has already been learned by A is harder than learning an unlearned pattern. Indeed, because of inhibition, A will prevent B from firing immediately after it, so the period available for B to fire is actually shorter than the length of the pattern. The problem is thus equivalent to that of finding a shorter pattern, which is obviously harder. This also means that neurons will effectively “prefer” to learn an unlearned pattern (and track back through it) than stacking on an almost saturated pattern. Again, this phenomenon will help distributing neurons harmoniously across patterns.

3 Discussion

As in our previous study, we have again found that STDP proves to be a robust mechanism for detecting and learning repeating spatiotemporal spike patterns in continuous spike trains (Masquelier et al., 2008). But by coupling it with a simple biologically plausible competitive mechanism, the studies presented here significantly extend our previous findings. First, we demonstrate that a given pattern can be divided in successive subpatterns, with one neuron coding for each subpattern, and that the neurons stack, learning in priority the first parts of the patterns. The subpattern durations should be of the order of the membrane time constant, which defines an optimal level of inhibition. Second, different patterns can be coded by different sets of (possibly overlapping) neurons. This powerful distributed coding scheme naturally emerges from both unsupervised STDP-based learning and lateral inhibition, which tends to prevent the neurons from learning the same subpatterns.

STDP had been shown to be capable of learning spatiotemporal patterns in both discrete spike volleys (Song, Miller, & Abbott, 2000; Gerstner & Kistler, 2002; Guyonneau, Van Rullen, & Thorpe, 2005; Masquelier & Thorpe, 2007) and periodic inputs (Gerstner et al., 1993; Gerstner, Kempter, van Hemmen, & Wagner, 1996; Yoshioka, 2002; Finelli et al., 2008; Hosaka, Araki, & Ikeguchi, 2008). In our previous study (Masquelier et al., 2008), we showed how STDP can detect a single repeating pattern in continuous




Figure 5: Latencies with multiple patterns. Here we show a typical result with three patterns (shown on each row) and nine neurons (in columns). We plotted the latencies for each neuron with respect to the beginning of each pattern. Again the latency is said to be zero if the neuron fired outside a pattern (i.e., generated a false alarm). For instance, it can be seen that neuron 1 became selective to pattern 3 and ends up with a latency of 19 ms, a hit rate of 99%, and no false alarms (from the point of view of patterns 1 and 2, it generates only false alarms). Pattern 1 has also been learned by neurons 5 and 8 with latencies of, respectively, 7 and 30 ms, with no false alarms and a hit rate of 100%. Neurons 3 and 4 learned pattern 1. Neurons 6 and 7 learned pattern 2. Neurons 2 and 9 “died” —stopped firing after too many (about 180) discharges outside the patterns, causing too much depression.

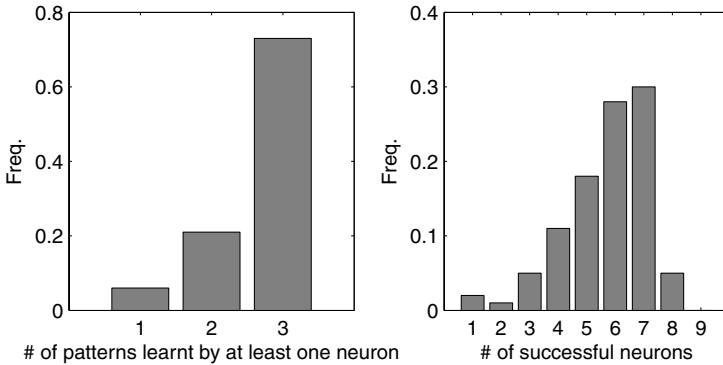


Figure 6: Performance with multiple patterns. The left panel plots the distribution of the number of patterns learned by at least one neuron across the 100 simulations. In more than two-thirds of the cases, all three patterns were learned by at least one neuron. The right panel plots the distribution of the number of successful neurons (criteria: hit rate above 90%, false alarm rate below 1 Hz). The mean is $5.71/9 = 63\%$.

spike trains and track back through it. That finding follows from the fact that STDP reinforces synapses that correspond to correlated spike trains (Kempster, Gerstner, & van Hemmen, 1999), and its sensitivity to the earliest predictors of subsequent spike events (Mehta, Quirk, & Wilson, 2000; Song et al., 2000). But to our knowledge, this is the first demonstration that by coupling it with a competitive mechanism, it can detect multiple spike patterns that occur at random intervals in continuous spike trains and to divide them in subpatterns. It is also the first time that a stacking effect has been predicted. Subsequent neurons with appropriate conduction delays could easily code for the whole patterns at low cost (Izhikevich, 2006). Specifically, the conduction delays should compensate for the differences of latencies, so that the subsequent neurons receive the input spikes simultaneously if and only if the subpatterns are presented in the correct order, that is, when the whole pattern is presented.

There is no consensus on the definition of a spike pattern, and we admit that ours is quite simple. Here, a pattern is seen as a succession of coincidences. The neurons become selective to successive coincidences of the patterns, whereas the exact temporal delays between the nearly simultaneous spikes of one coincidence do not matter. Note, however, that the succession of the coincidences is another temporal aspect, which, even though not relevant after convergence, greatly facilitates the learning phase and allows the neurons to stack. It would be much harder to find single repeating coincidences on their own. But because the much longer sequences repeat (here lasting 50 ms), the neurons only need to fire somewhere within a

50 ms window to start learning the patterns and track back through them until stacking from the beginning of them.

We emphasize that the proposed learning scheme is completely unsupervised. No teaching signal tells the neuron when to learn. Nor is there any mechanism for labeling the input patterns. Biologically plausible mechanisms for supervised learning of spike patterns have also been proposed (Gütig & Sompolinsky, 2006), sometimes using STDP in addition (Legenstein, Naeger, & Maass, 2005).

It is known that in feedforward networks, no computational power is lost if, like here, inhibition is used exclusively for unspecific lateral inhibition, and no flexibility is lost if, like here, plasticity is restricted to excitatory synapses (Maass, 2000). The circuit we propose is thus both powerful and in accordance with the asymmetry between inhibition and excitation observed in the cortex: only about 15% of synapses are inhibitory, evidence for their direct involvement in plasticity is somewhat controversial, and their action is usually found to be local and nonselective.

The classic *k*-winner-take-all (kWTA) problem has generally been studied with rate-based coding frameworks in which it means that the most frequently firing neurons should inhibit their competitors and thereby stop them from firing. Biophysical circuits to perform such operations (and the particular case of 1WTA, i.e., max operation) have been proposed (Elias & Grossberg, 1975; Amari & Arbib, 1977; Yuille & Grzywacz, 1989; Coultrip, Granger, & Lynch, 1992; Yu, Giese, & Poggio, 2002; Knoblich, Bouvrie, & Poggio, 2007; Kouh & Poggio, 2008). However, it is important to realize that these circuits are significantly more complex than what we propose, and they typically take several tens of milliseconds to reach a steady regime. The duration of the transient increases if input values are close to each other (Knoblich et al., 2007). With our approach, in which the latencies of single spikes are reliable (as opposed to stochastic, e.g., Poissonian) and carry information, kWTA simply means that the earliest firing neurons should prevent their competitors from firing for a while—something that lateral inhibition can easily do (Thorpe, 1990).

The mechanism presented here is generic and could be at work in all brain areas. Nevertheless, if we focus on sensory systems, and in particular audition or vision, it is easy to imagine how competitive STDP learning could generate a vast dictionary of features from experience. The model predicts that frequently occurring features are not only more likely to be learned, but also will be processed and recognized faster than unfamiliar ones (recall that postsynaptic latencies decrease with training). This last point is in accordance with psychophysical results showing that familiar categories such as faces are processed faster (Thorpe, Crouzet, Kirchner, & Fabre-Thorpe, 2006; Crouzet, Thorpe, & Kirchner, 2007) or that processing times can be speeded up with experience (Masquelier, 2008).

In the brain, evidence for reliable spike timing is growing (see Tiesinga, Fellous, & Sejnowski, 2008, for a recent review), and electrophysiologists

have reported the existence of repeating spatiotemporal spike patterns with millisecond precision, both *in vitro* and *in vivo*. These repeating patterns can last from a few tens of ms to several seconds (Frostig, Frostig, & Harper, 1990; Prut et al., 1998; Fellous, Tiesinga, Thomas, & Sejnowski, 2004; Ikegaya et al., 2004; Abeles, 2004; Luczak, Barthó, Marguet, Buzsáki, & Harris, 2007; Rolston, Wagenaar, & Potter, 2007). In this letter, we propose a cheap and robust mechanism that can detect, learn, and represent these repeating spatiotemporal patterns. The hypotheses we made (essentially the presence of LIF neurons, STDP, and lateral inhibition based kWTA) are surprisingly weak. The results thus constitute a strong argument in favor of temporal coding, that is, the hypothesis that the brain uses the spike times, and not only the firing rates, to encode, decode, and process information.

Appendix: Code

The simulations were performed using Matlab R14. The source code is available from the authors on request.

A.1 Spike Trains. The methods used for generating the spike trains were essentially identical to the ones used in our earlier study (Masquelier et al., 2008), except that multiple patterns were copy-pasted. They are described again here for the reader's convenience.

The spike trains were prepared before the simulation (Figure 1 illustrates the type of spike trains we used, though with a smaller set of neurons). Because of memory constraints, we did not have a single nonrepeating sequence. Instead, a spike sequence lasting one-third of the total simulation time was repeated three times per run, but this repetition had no impact on the results. Each afferent emits spikes independently using a Poisson process with an instantaneous firing rate r that varies randomly between 0 and 90 Hz. The maximal rate change s was chosen so that the neuron could go from 0 to 90 Hz in 50 ms. Time was discretized using a time step dt of 1 ms. At each time step:

1. The afferent has a probability of $r \cdot dt$ of emitting a spike (whose exact date is then picked randomly within the 1 ms time bin).
2. Its instantaneous firing rate is modified: $dr = s \cdot dt$, where s is the speed of rate change (in Hz/s), and clipped in $[0, 90]$ Hz.
3. Its speed of rate change is modified by ds , randomly picked from a uniform distribution over $[-360 + 360]$ Hz/s, and clipped in $[-1800 + 1800]$ Hz/s.

We chose to apply the random change to s as opposed to r so as to have a continuous s function and a smoother r function.

A limitation of this work is the excitatory-only scheme. Consequently, something like "afferent A must not spike" cannot be learned, and learning will be limited to positive patterns. We thus wanted repeating patterns in

which all the afferents spike at least once. We could have made up such patterns, but we wanted the pattern to have exactly the same statistics as the Poisson distracter part (to make the pattern detection harder), so we preferred to randomly pick 50 ms periods from the original Poisson spike trains and then to copy-paste them (see below). To make sure these randomly selected periods all contain at least one spike from each afferent, we implemented a mechanism that triggers a spike whenever an afferent has been silent for more than 50 ms (leading to a minimal firing rate of 20 Hz). Clearly, such a mechanism is not implemented in the brain; it is just an artifice we used here to make the pattern detection harder. As a result, the average firing rate was 54 Hz, and not the 45 Hz we would have without this additional mechanism.

Once the random spike train had been generated, some randomly selected portions were selected and used to define the patterns to be repeated, and these were copy-pasted at irregular intervals throughout the sequence. For each pattern, the copy-paste did not involve all the afferents: half of them (randomly picked) conserved their original spike trains. But we discretized the spike trains of the other half of afferents into 50 ms sections. We randomly picked one of these sections and copied the corresponding spikes. Then we randomly picked a certain number of these sections ($1/3 \times 1/\text{number of patterns}$, so that the total time with pasted patterns represented one-third of the simulation time), avoiding consecutive ones, and replaced the original spikes by the copied ones. Jitter was added before the pasting operation, picked from a gaussian distribution with a mean of zero and standard deviation 1 ms. We repeated the copy-paste operation as many times as the number of patterns we wanted, avoiding the situation where two different patterns are pasted within the same 50 ms period. For each pattern, the involved afferents were independently randomly chosen; the selected sets were thus not disjoint.

After this copy-paste operation, a 10 Hz Poissonian spontaneous activity was added to all neurons and all the time. The total activity was thus 64 Hz on average, and spontaneous activity represented about 16% of it.

A.2 Leaky Integrate-and-Fire Neuron. The neuronal model is the same as the one we used in our earlier study (Masquelier et al., 2008), except that inhibitory postsynaptic potentials (IPSPs) are now involved.

Specifically, we modeled the LIF neuron using Gerstner's spike response model (SRM) (Gerstner et al., 1993). That is, instead of solving the membrane potential differential equation, we used kernels to model the effect of spikes on the membrane potential.

Each presynaptic spike j , with arrival time t_j , is supposed to add to the membrane potential an excitatory postsynaptic potential (EPSP) of the form

$$\varepsilon(t - t_j) = K \cdot \left(\exp\left(-\frac{t - t_j}{\tau_m}\right) - \exp\left(-\frac{t - t_j}{\tau_s}\right) \right) \cdot \Theta(t - t_j),$$

where τ_m is the membrane time constant (here 10 ms), τ_s is the synapse time constant (here 2.5 ms), and Θ is the Heaviside step function:

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

and K is just a multiplicative constant chosen so that the maximum value of the kernel is 1 (the voltage scale is arbitrary in this letter).

The last emitted postsynaptic spike i has an effect on the membrane potential modeled as follows:

$$\eta(t - t_i) = T \cdot \left(K_1 \cdot \exp\left(-\frac{t - t_i}{\tau_m}\right) - K_2 \cdot \left(\exp\left(-\frac{t - t_i}{\tau_m}\right) - \exp\left(-\frac{t - t_i}{\tau_s}\right) \right) \right) \cdot \Theta(t - t_i),$$

where T is the threshold of the neuron (here 550, arbitrary units). The first term models the positive pulse and the second one the negative spike afterpotential that follows the pulse. Here we used $K_1 = 2$ and $K_2 = 4$.

When a neuron fires at time t_k , it sends to the others an inhibitory postsynaptic potential (IPSP). For simplicity, we used the same kernel as for EPSP with a multiplicative constant α (0.25 in the baseline condition) used to tune the inhibition strength:

$$\mu(t - t_k) = -\alpha \cdot T \cdot \varepsilon(t - t_k).$$

ε , η , and μ kernels were rounded to zero when, respectively, $t - t_j$, $t - t_i$, and $t - t_k$ were greater than $7\tau_m$.

At any time, the membrane potential is

$$p = \eta(t - t_i) + \sum_{j/t_j > t_i} w_j \cdot \varepsilon(t - t_j) + \sum_{k/t_k > t_i} \mu(t - t_k),$$

where the w_j are the excitatory synaptic weights, between 0 and 1 (arbitrary units).

This SRM formulation allows us to use event-driven programming: we compute the potential only when a new presynaptic spike is integrated. We then estimate numerically if the corresponding EPSP will cause the threshold to be reached in the future and at what date. If it is the case, a postsynaptic spike is scheduled. Such postsynaptic spike events cause all the EPSPs and IPSPs to be flushed, and a new t_i is used for the η kernel. There is then a refractory period of 5 ms when the neuron is not allowed to fire.

A.3 Spike-Timing-Dependent Plasticity. The STDP model we used is identical to the one used in our earlier study (Masquelier et al., 2008). It is described again here for the reader's convenience.

We used an additive exponential update rule:

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

Following learning, the weights were clipped to [0,1]. Note that such a bounded additive rule leads to a bimodal distribution after convergence (VanRossum, Bi, & Turrigiano, 2000).

We used $\tau^+ = 16.8$ ms and $\tau^- = 33.7$ ms, in accordance with experimental measurements (Bi & Poo, 2001). We restricted the learning window to $[t_i - 7 \cdot \tau^+, t_i]$ for LTP and to $[t_i, t_i + 7 \cdot \tau^-]$ for LTD. For each afferent, we also limited LTP (resp. LTD) to the last (first) presynaptic spike before (after) the postsynaptic one (nearest spike approximation). We did not take the effects of finer triplets of spikes (Pfister & Gerstner, 2006) into account.

It was found that small learning rates led to more robust learning. We used $a^+ = 0.03125$ and $a^- = 0.85 \cdot a^+$.

Importantly, $a^- \tau^- > a^+ \tau^+$, meaning STDP function is biased towards depression (Song et al., 2000). This means, among other things, that if no repeating patterns were inserted, STDP would thus gradually decrease the synaptic weights until the threshold would no longer be reached.

Acknowledgments

This research was supported by the CNRS, STREP Decisions-in-Motion (IST-027198), the Agence National de Recherche (Projects Natstats and Hearing in Time), and SpikeNet Technology SARL.

References

- Abeles, M. (2004). Neuroscience: Time is precious. *Science*, 304(5670), 523–524.
- Amari, S., & Arbib, M. (1977). *Systems neuroscience*. San Diego, CA: Academic Press.
- Bi, G., & Poo, M. (1998). Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.*, 18(24), 10464–10472.
- Bi, G., & Poo, M. M. (2001). Synaptic modification by correlated activity: Hebb's postulate revisited. *Ann. Rev. Neurosci.*, 24, 139–166.
- Cassenaar, S., & Laurent, G. (2007). Hebbian STDP in mushroom bodies facilitates the synchronous flow of olfactory information in locusts. *Nature*, 448(7154), 709–713.
- Coultrip, R., Granger, R., & Lynch, G. (1992). A cortical model of winner-take-all competition via lateral inhibition. *Neural Networks*, 5, 47–54.

- Crouzet, S., Thorpe, S., & Kirchner, H. (2007). Category-dependent variations in visual processing time. *Journal of Vision*, 7, 922a.
- Delorme, A., Perrinet, L., Thorpe, S., & 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.
- Elias, S., & Grossberg, S. (1975). Pattern formation, contrast control, and oscillations in the short term memory of shunting on-center off-surround networks. *Biol. Cybern.*, 20, 69–98.
- Feldman, D. (2000). Timing-based LTP and LTD at vertical inputs to layer II/III pyramidal cells in rat barrel cortex. *Neuron*, 27(1), 45–56.
- Fellous, J. M., Tiesinga, P. H., Thomas, P. J., & Sejnowski, T. J. (2004). Discovering spike patterns in neuronal responses. *J. Neurosci.*, 24(12), 2989–3001.
- Finelli, L. A., Haney, S., Bazhenov, M., Stopfer, M., & Sejnowski, T. J. (2008). Synaptic learning rules and sparse coding in a model sensory system. *PLoS Comput. Biol.*, 4(4), e1000062.
- Frostig, R. D., Frostig, Z., & Harper, R. M. (1990). Recurring discharge patterns in multiple spike trains. *Biol. Cybern.*, 62(6), 487–493.
- Gerstner, W., Kempter, R., van Hemmen, J. L., & Wagner, H. (1996). A neuronal learning rule for sub-millisecond temporal coding. *Nature*, 383(6595), 76–81.
- Gerstner, W., & Kistler, W. (2002). *Spiking neuron models*. Cambridge: Cambridge University Press.
- Gerstner, W., Ritz, R., & van Hemmen, J. L. (1993). Why spikes? Hebbian learning and retrieval of time-resolved excitation patterns. *Biol. Cybern.*, 69(5–6), 503–515.
- Gütig, R., & Sompolinsky, H. (2006). The tempotron: A neuron that learns spike timing-based decisions. *Nat. Neurosci.*, 9(3), 420–428.
- Guyonneau, R., VanRullen, R., & Thorpe, S. (2004). Temporal codes and sparse representations: A key to understanding rapid processing in the visual system. *J. Physiol. Paris.*, 98(4–6), 487–497.
- Guyonneau, R., VanRullen, R., & Thorpe, S. (2005). Neurons tune to the earliest spikes through STDP. *Neural Comput.*, 17(4), 859–879.
- Hosaka, R., Araki, O., & Ikeguchi, T. (2008). STDP provides the substrate for igniting synfire chains by spatiotemporal input patterns. *Neural Comput.*, 20(2), 415–435.
- Ikegaya, Y., Aaron, G., Cossart, R., Aronov, D., Lampl, I., Ferster, D., et al. (2004). Synfire chains and cortical songs: Temporal modules of cortical activity. *Science*, 304(5670), 559–564.
- Izhikevich, E. M. (2006). Polychronization: Computation with spikes. *Neural Comput.*, 18(2), 245–282.
- Jacob, V., Brasier, D. J., Erchova, I., Feldman, D., & Shulz, D. E. (2007). Spike timing-dependent synaptic depression in the in vivo barrel cortex of the rat. *J. Neurosci.*, 27(6), 1271–1284.
- Kempter, R., Gerstner, W., & van Hemmen, J. L. (1999). Hebbian learning and spiking neurons. *Phys. Rev. E*, 59(4), 4498–4514.
- Knoblich, U., Bouvrie, J., & Poggio, T. (2007). *Biophysical models of neural computation: Max and tuning circuits*. (Tech. Report). Cambridge, MA: MIT.
- Kouh, M., & Poggio, T. (2008). A canonical neural circuit for cortical nonlinear operations. *Neural Comput.*, 20(6), 1427–1451.

- Legenstein, R., Naeger, C., & Maass, W. (2005). What can a neuron learn with spike-timing-dependent plasticity? *Neural Comput.*, 17(11), 2337–2382.
- Luczak, A., Barthó, P., Marguet, S. L., Buzsáki, G., & Harris, K. D. (2007). Sequential structure of neocortical spontaneous activity in vivo. *Proc. Natl. Acad. Sci. USA*, 104(1), 347–352.
- Maass, W. (2000). On the computational power of winner-take-all. *Neural Comput.*, 12(11), 2519–2535.
- Markram, H., Lübke, J., Frotscher, M., & Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APS and EPSPS. *Science*, 275(5297), 213–215.
- Masquelier, T. (2008). *Learning mechanisms to account for the speed, selectivity and invariance of responses in the visual cortex*. Unpublished doctoral dissertation, Université Paul Sabatier Toulouse 3.
- Masquelier, T., Guyonneau, R., & Thorpe, S. J. (2008). Spike timing dependent plasticity finds the start of repeating patterns in continuous spike trains. *PLoS ONE*, 3(1), e1377.
- Masquelier, T., & Thorpe, S. J. (2007). Unsupervised learning of visual features through spike timing dependent plasticity. *PLoS Comput. Biol.*, 3(2), e31.
- Mehta, M. R., Quirk, M. C., & Wilson, M. A. (2000). Experience-dependent asymmetric shape of hippocampal receptive fields [see comments]. *Neuron*, 25(3), 707–715.
- Meliza, C. D., & Dan, Y. (2006). Receptive-field modification in rat visual cortex induced by paired visual stimulation and single-cell spiking. *Neuron*, 49(2), 183–189.
- Mu, Y., & Poo, M. M. (2006). Spike timing-dependent LTP/LTD mediates visual experience-dependent plasticity in a developing retinotectal system. *Neuron*, 50(1), 115–125.
- Pfister, J., & Gerstner, W. (2006). Triplets of spikes in a model of spike timing-dependent plasticity. *Journal of Neuroscience*, 26(38), 9673–9682.
- Prut, Y., Vaadia, E., Bergman, H., Haalman, I., Sloviter, H., & Abeles, M. (1998). Spatiotemporal structure of cortical activity: Properties and behavioral relevance. *J. Neurophysiol.*, 79(6), 2857–2874.
- Rolston, J. D., Wagenaar, D. A., & Potter, S. M. (2007). Precisely timed spatiotemporal patterns of neural activity in dissociated cortical cultures. *Neuroscience*, 148(1), 294–303.
- Song, S., Miller, K., & Abbott, L. (2000). Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nat. Neurosci.*, 3(9), 919–926.
- Thorpe, S. (1990). Spike arrival times: A highly efficient coding scheme for neural networks. In R. Eckmiller, G. Hartmann, & G. Hauske (Eds.), *Parallel processing in neural systems and computers* (pp. 91–94). Dordrecht: Elsevier.
- Thorpe, S., Crouzet, S., Kirchner, H., & Fabre-Thorpe, M. (2006). Ultra-rapid face detection in natural images: Implications for computation in the visual system. In *Proceedings of the First French Conference on Computational Neurosciences* (pp. 124–127). Ponta Mousson, France.
- Tiesinga, P., Fellous, J.-M., & Sejnowski, T. J. (2008). Regulation of spike timing in visual cortical circuits. *Nat. Rev. Neurosci.*, 9(2), 97–107.
- VanRossum, M., Bi, G., & Turrigiano, G. (2000). Stable Hebbian learning from spike timing-dependent plasticity. *Journal of Neuroscience*, 20(23), 8812–8821.

- Vislay-Meltzer, R. L., Kampff, A. R., & Engert, F. (2006). Spatiotemporal specificity of neuronal activity directs the modification of receptive fields in the developing retinotectal system. *Neuron*, *50*(1), 101–114.
- Yoshioka, M. (2002). Spike-timing-dependent learning rule to encode spatiotemporal patterns in a network of spiking neurons. *Phys. Rev. E. Stat. Nonlin. Soft. Matter Phys.*, *65*(1 Pt. 1), 011903.
- Young, J. M., Waleszczyk, W. J., Wang, C., Calford, M. B., Dreher, B., & Obermayer, K. (2007). Cortical reorganization consistent with spike timing—but not correlation-dependent plasticity. *Nat. Neurosc.*, *10*(7), 887–895.
- Yu, A. J., Giese, M. A., & Poggio, T. (2002). Biophysiologicaly plausible implementations of the maximum operation. *Neural Comp.*, *14*(12), 2857–2881.
- Yuille, A. L., & Grzywacz, N. M. (1989). A winner-take-all mechanism based on presynaptic inhibition feedback. *Neural Comp.*, *1*(3), 334–347.
- Zhang, L., Tao, H., Holt, C., Harris, W., & Poo, M. (1998). A critical window for cooperation and competition among developing retinotectal synapses. *Nature*, *395*(6697), 37–44.

Received June 16, 2008; accepted September 23, 2008.