LETTER **Communicated by Laurence Abbott**

Neurons Tune to the Earliest Spikes Through STDP

Rudy Guyonneau
rudy.guyonneau@cerco.ups-tlse.fr
Centre de Recherche "Cerveau et Cognition," Toulouse 31000, France, and Spikenet Technology, Revel, France

Rufin VanRullen
Rufin.vanrullen@cerco.ups-tlse.fr
Centre de Recherche "Cerveau et Cognition," Toulouse 31000, France

Simon J. Thorpe
Simon.Thorpe@cerco.ups-tlse.fr
Centre de Recherche "Cerveau et Cognition, Toulouse 31000, France, and Spikenet Technology, Revel 31250, France

Spike timing-dependent plasticity (STDP) is a learning rule that modifies the strength of a neuron's synapses as a function of the precise temporal relations between input and output spikes. In many brains areas, temporal aspects of spike trains have been found to be highly reproducible. How will STDP affect a neuron's behavior when it is repeatedly presented with the same input spike pattern? We show in this theoretical study that repeated inputs systematically lead to a shaping of the neuron's selectivity, emphasizing its very first input spikes, while steadily decreasing the postsynaptic response latency. This was obtained under various conditions of background noise, and even under conditions where spiking latencies and firing rates, or synchrony, provided conflicting informations. The key role of first spikes demonstrated here provides further support for models using a single wave of spikes to implement rapid neural processing.

1 Introduction

Activity-dependent learning at the systems level relies on dynamical modifications of synaptic strength at the cellular level. Experimental data have shown that these modifications depend on temporal pairing between a pre- and a postsynaptic spike: an excitatory synapse receiving a spike before a postsynaptic one is emitted potentiates while it weakens the other way around (Markram, Lubke, Frotscher, & Sakmann, 1997). The amount of modification depends on the delay between these two events: maximal when pre- and postsynaptic spikes are close together, the effects gradually decrease

Neural Computation **17, 859–879** (2005) © 2005 Massachusetts Institute of Technology

and disappear with intervals in excess of a few tens of milliseconds (Bi & Poo, 1998; Zhang, Tao, Holt, Harris, & Poo, 1998; Feldman, 2000). Learning at the neuronal level thus heavily depends on the temporal structure of neuronal responses.

Interestingly, in many brain areas, the temporal precision of spikes during stimulus-locked responses can be in the millisecond range. The evidence is clear in both the auditory cortex (Heil, 1997) and the somatosensory system (Petersen, Panzeri, & Diamond, 2001). Reproducible temporal structure can also be found in the visual system, such as in MT (Bair & Koch, 1996), and from retinal ganglion cells to the inferotemporal cortex (Sestokas, Lehmkuhle, & Kratz, 1991; Meister & Berry, 1999; Liu, Tzonev, Rebrik, & Miller, 2001; Richmond & Optican, 1990; Victor & Purpura, 1996; Nakamura, 1998).

This raises the question of how a neuron with STDP will respond when repeatedly stimulated with the same pattern of input spikes. Let us consider a simple case where a neuron is exposed to an input spike pattern

---

Figure 1: *Facing page*. Single wave experiment. For a repeated spike wave under realistic conditions, a neuron learns to react faster to its target. Synaptic weights converge onto the earliest firing afferents, even with 25 ms jitter or 50 Hz background activity. (A) Typical incoming activity. Bottom: raster plot of a jittered spike wave amid spontaneous activity. At each presentation, 5 ms jitter and 5 Hz background activity are regenerated; the resulting pattern is presented to the postsynaptic neuron (with initial potential set to 0); when it spikes, the STDP learning rule is applied and its potential reset to 0 before going to the next presentation. Prior to presentation, the presynaptic neurons do not fire any spike. Top: the corresponding poststimulus time histogram (PSTH): the gaussian form in the left-most part corresponds to the reproduced spike wave. (B) Dynamics of repeated STDP. Top: sum of all synaptic weights at each presentation (the dashed line represents the output neuron threshold). The sum of the synaptic weights stored in the afferents stabilizes at threshold value. Bottom: the horizontal axis corresponds to the number of presentations (i.e., the learning step). The black line refers to the left axis and shows the reduction of postsynaptic latency during the course of learning. The background image refers to the right axis where each synapse weight is mapped by a gray-level index (see the corresponding bar on the right). Synapses are ordered by spiking latency of the corresponding neuron within the original reproducible input pattern (here, a wave), more precisely before superimposing spontaneous activity or adding some jitter. This order is decided a priori and stays fixed for the whole simulation. During learning, earliest synapses become fully potentiated and later ones are weakened. (C) Effect of jitter. Jitter is generated by a gaussian distribution. Increasing its standard deviation does not affect convergence until about 10 ms. From there, it slows the system roughly quadratically. (D) Effect of spontaneous firing rate. Increasing background activity slows convergence approximatively linearly.

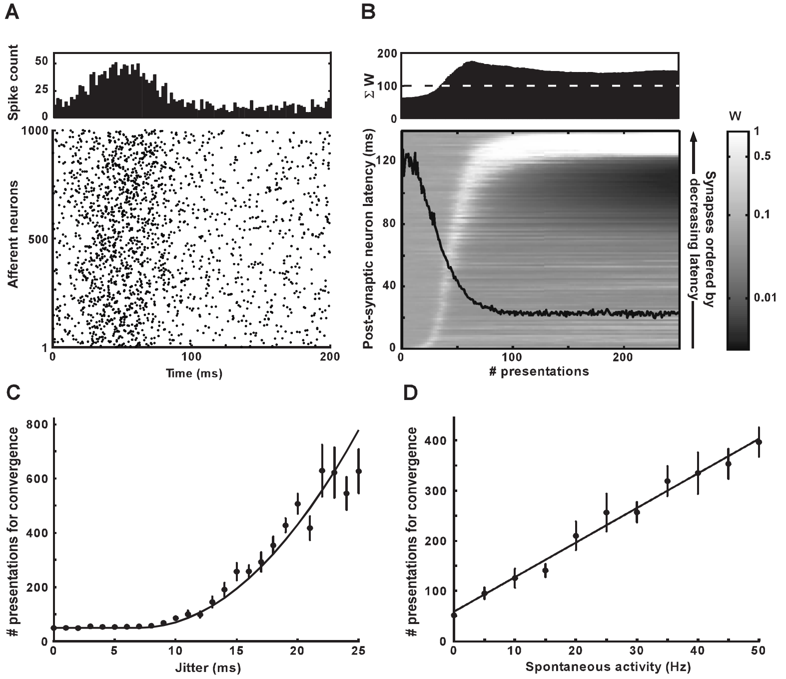

consisting of one input spike on each afferent synapse, with arbitrary but fixed (e.g., gaussian distributed, as in latencies. For one given input pattern presentation, the input spikes elicit a postsynaptic response, triggering the STDP rule. Synapses carrying input spikes just preceding the postsynaptic one are potentiated, while later ones are weakened. The next time this input pattern is re-presented, firing threshold will be reached sooner, which implies a slight decrease of the postsynaptic spike latency. Consequently, the learning process, while depressing some synapses it had previously potentiated, will now reinforce different synapses carrying even earlier spikes than the preceding time. By iteration, it follows that when the presentation of the same input spike pattern is repeated, the postsynaptic spike latency will tend to stabilize at a minimal value while the first synapses become fully potentiated and later ones fully depressed (see Figure 4 in Song, et al., 2000, as well as Gerstner and Kistler, 2002a, for related demonstrations; see also Figures 1A and 1B).

Under these simplistic conditions, neurons can learn to react faster to a given stimulus pattern by emphasizing the role of the earliest firing inputs. However, input spike patterns in the brain do not consist of one-spike-per-neuron waves of infinite precision. Not only do stimulus-locked responses occur with a jitter on the order of 1 or more milliseconds (Mainen

& Sejnowski, 1995), neocortical cells also spontaneously fire at rates up to 20 Hz in awake animals (Evarts, 1964; Hubel, 1959; Steriade, Oakson, & Kitsikis, 1978). Finally, temporal structure in spike trains is often distributed over numerous consecutive spikes in long time windows. With realistic conditions, would the effect of STDP still emphasize the very first afferents? Would this be the case even when other types of neural codes (firing rate, synchrony) are present within the spike trains? This would have critical implications for our understanding of neural information coding and processing. In this theoretical study, we seek to answer these questions through a range of simulations using biologically plausible neuron models.

2 Spike Waves

A target neuron is repeatedly presented with an input spike pattern. In a first example, it consists of a single asynchronous spike wave: one spike for each synapse with gaussian-distributed latencies (50 ms mean, 20 ms width). At each presentation, spike times are jittered, and Poisson spontaneous activity is added to the spike trains. We varied the amount of jitter and spontaneous firing and investigated their effect on the target neuron learning behavior.

2.1 Methods.

2.1.1 Integrate and Fire. The postsynaptic neuron is connected to a thousand afferent neurons via as many synapses and integrates spikes across time (in all simulations except one, no leakage was involved). Systematically starting at a resting level of 0 before presentation, it sums the weight of the synapses that carry each incoming spike. The presynaptic neurons are assumed to be quiet prior to presentation; they do not fire any spike. When reaching its threshold (100 for all simulations), the postsynaptic neuron fires, triggers an STDP-inspired learning rule for its first action potential only, and then sets its potential back to a resting level of 0.

2.1.2 STDP Model. The learning function F has the prototypical form of STDP according to electrophysiological results in cultured hippocampal neurons (Bi & Poo, 1998). The amount of synaptic modification arising from a single pair of pre- and postsynaptic spikes separated by time Δt is expressed as follows:

$$\begin{array}{lll} \text{if} & \tau_{+} \leq \Delta t \leq 0, & F(\Delta t) = A_{+}.\,(1 - (\Delta t/\tau +)) \\ \text{if} & 0 < \Delta t \leq \tau_{-}, & F(\Delta t) = -A_{-}.\,(1 - (\Delta t/\tau -)) \\ \text{otherwise} & & F(\Delta t) = 0, \end{array}$$

where A_{+} and A_{-} determine the maximum amounts of synaptic modification that occur when Δt is close to zero ($A_{+} = A_{-} = 1$ in all simulations). τ_{+} and τ_{-} are the temporal windows for, respectively, potentiation and depression, expressed in milliseconds (here, $\tau_+ = -20$ ms and $\tau_- = 22$ ms). Synaptic growth from learning step n to $n+1$ is computed as follows:

$$w_i(n+1) = w_i(n) + F(t_{\mathrm{pre}} - t_{\mathrm{post}}),$$

where $w_i(n)$ is the "free" weight of synapse i at nth presentation step, t_{post} the postsynaptic spike timing, and t_{pre} the presynaptic spike timing. Under these conditions, synaptic weights may grow to infinitely large values. To address this obstacle, we implemented a sigmoidal saturation function g:

$$W_i = g(w_i),$$

where $g(x) = ((\pi/2) + \mathrm{atan}(x))/\pi$. The "free" weight, labeled w_i, corresponds to the unconstrained weight of the synapse, which connects input neuron$i$ to the postsynaptic one. It ranges from $-\infty$ to $+\infty$ and is the target of the strengthening/weakening process. The "effective" weights, W_i, lie in the $]w_{\mathrm{min}}, w_{\mathrm{max}}[$ interval (here $w_{\mathrm{min}} = 0, w_{\mathrm{max}} = 1$). These "effective" weights are the ones that are considered for calculating the excitatory postsynaptic potential.[1]

2.1.3 Jitter. When needed, each reproducible spike was displaced in time by a value chosen randomly at each presentation from a gaussian distribution of mean 0 and standard deviation set to the expected jitter.

2.1.4 Spontaneous Activity. Spontaneous activity is generated using a Poisson process as described in section 3.1. It is redrawn for each afferent neuron, at each presentation, before being superimposed on the jittered reproducible structure to constitute the incoming activity.

2.1.5 Convergence Criterion. Convergence is met when the local average of postsynaptic latencies, in a ±5-step window, remains within 1 ms for 100 consecutive steps (reported convergence is the first of these 100 steps). We also checked that the synaptic weights were indeed selected on the basis of their earliest inputs.

---

[1] Note that here, contrary to other update rules (Rubin, Lee, & Sompolinsky, 2001), when the effective weight approaches the upper and lower bounds, both the expected amounts of reward and punishment (in the "effective" weight domain) tend toward 0. In the multiplicative update case (Song et al., 2000), these amounts stay balanced one compared to the other: when the weight reaches a bound, for example, w_{max}, then reward goes to 0, whereas punition is maximal and inversely. For the additive update rule (Cateau & Fukai, 2003), both potentiation and depression are independent of the weight. If the update results in a synaptic weight outside the bounds, the weight is clipped to the boundary values.

2.2 Resistance to Jitter and Spontaneous Activity. A typical example with 5 ms jitter and 5 Hz spontaneous activity (see Figure 1A) shows that the target neuron learns to react faster to the input pattern—in the case here, a spike wave—by selectively reinforcing the earliest firing afferents (see Figure 1B). Initially, synaptic weights are set so as to evoke the first postsynaptic response when the entire reproducible pattern, the spike wave, has been integrated (in the present case, around 120 ms after stimulus onset). Synapses are then progressively reinforced and then weakened, shortening the postsynaptic neuron latency from one step to the following. This is especially visible between steps 20 and 100, where this codependent dynamic appears as a "bright crest" going from the latest to the earliest afferences. When the postsynaptic neuron latency stabilizes, the STDP rule is systematically applied at approximately the same time: its rewarding part constantly affects the same synapses, those receiving the earliest of the reproducible spikes, thus driving them to their maximal strengths. Symmetrically, later spikes invariably arrive in the weakening part of the STDP window: the corresponding synapses are continuously depressed.

This trend can also be observed in the evolution of the total amount of synaptic weight of the neuron at a given step (see Figure 1B, top). It corresponds to the potential the model neuron would reach if each of its synapses was hit by one spike only. Initially, it starts from a value well below the threshold needed for the postsynaptic neuron to fire a spike. Here, the presence of spontaneous activity allows the neuron to reach its threshold and trigger the STDP rule for the first time. As the learning process goes on, the sum of synaptic weights increases due to the increasing number of spikes, thus of synapses, falling in the rewarded part of the STDP window; tends toward the neuron's threshold value; and stabilizes around this level. Since the maximum synaptic weight is fixed at 1.0, the neuron can reach a state where the first N spikes suffice to make it fire (where N is the output neuron's threshold).

The neuron will fire early only if it receives more or less the same pattern as the one learned during the first moments of incoming activity. In fact, it responds increasingly faster to a precise sequence of spikes than to any other (see Figure 5).

Simulations showed that the convergent behavior of this trend is very robust. When jitter is increased (with spontaneous activity set to 0), the number of presentations needed for convergence stays the same until the amount of jitter reaches 10 ms; then it increases roughly quadratically (see Figure 1C). Thus, jitter has little or no effect on the consequences of repeated STDP for plausible values. Moreover, convergence of the synaptic weights onto the earliest afferents could even be obtained with jitter in the range of 20 to 25 ms, that is, as wide as the input spike times distribution itself.

When increasing spontaneous activity (in the 5–50 Hz range without any jitter), the convergence is also slowed but is nonetheless obtained even at the highest rate (see Figure 1D). Note that in this case, the input spike wave represents less than 15% of the total spikes. As only the 10% of synapses carrying the earliest spikes are selected, this result implies that the STDP rule is able to focus exclusively on less than 2% of the spikes, those whose timing is reproducible and early, while discarding the rest.

3 Spike Trains

When a single wave of spikes is the reproducible structure, it is clear that convergence to a state in which weights are heavily concentrated on the earliest firing inputs is very robust. However, it could be argued that the STDP rule focuses on the left-most part of the train only because it is the only part that contains information from one pattern presentation to the next (indeed, the right-most part contains only randomly generated spikes).

3.1 Uniform Spike Distribution. The main simulation uses an input spike pattern in the form of 500 ms–long spike trains where the temporal structure and the amount of information is statistically homogeneous across time and across afferent neurons. More specifically, spike trains were generated according to a Poisson process. Each of the 1000 excitatory afferents emits a given spike train where each interspike interval, *isi*, was determined according to a Poisson rule process depending on the expected rate of the train, u:

$$isi(r) = \ln(r)/-u,$$

where r is a random value chosen from a uniform distribution on the interval]0.0, 1.0[. The reproducible structure for the simulation illustrated in Figure 2 is defined once and for all using this method ($u = 20$ Hz for all afferent neurons). Note that we are not suggesting that neuronal responses are completely stochastic, since reproducibility implies the contrary (see, e.g., Meister & Berry, 1999). We need this only to ensure that no a priori assumptions are made as to how neural information is encoded, thanks to the homogeneous nature of Poisson processes. This basic pattern is reproduced on each presentation, undergoing 5 ms jitter and mixed with 5 Hz spontaneous activity (see Figure 2A).

The initial synaptic weights are set so as to elicit the first postsynaptic response after approximately 400 ms. As in the previous simulation, latency decreases steadily from this initial value to stabilize after 1500 presentations. Conjointly, synapses carrying the first spikes become fully potentiated and later ones are fully depressed (see Figure 2B, bottom). The sum of synaptic weights displays the same dynamic as in the spike wave experiment (see section 2.2): going from a very low level, it rises until it reaches the minimal value needed to make the postsynaptic neuron fire on a single volley of

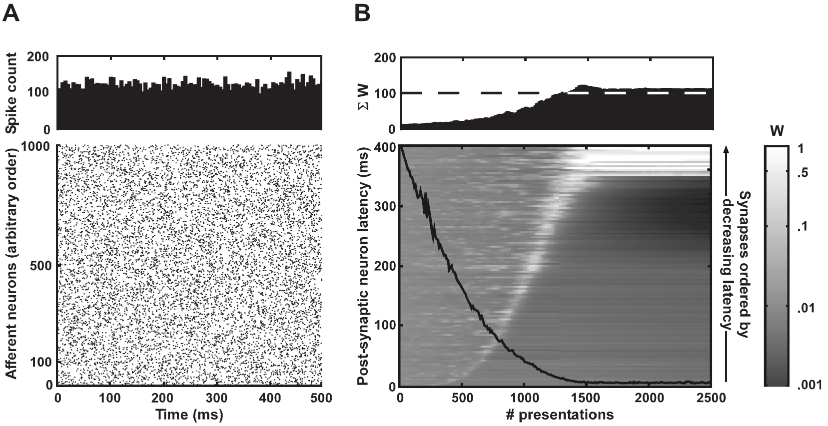

Figure 2: Spike trains experiment. (A) Typical incoming activity. The input spike pattern consists of arbitrary 500 ms–long spike trains, where the amount of information is equally distributed over time. At each presentation, this reproducible structure is modified using 5 ms jitter and 5 Hz spontaneous activity, then presented to the postsynaptic neuron, whose potential is set to 0. Prior to presentation, the presynaptic neurons do not fire any spike. Due to the random aspect of the input design (Poisson-inspired process; see section 3.1), latencies and firing rates in the reproducible input structure were quite variable, ranging, respectively, from 0 to ≈350 ms (mean of ≈52 ± 51 ms, standard deviation) and from 4 to 44 Hz (mean 19.6 ± 6.3 Hz, standard deviation). (B) Dynamics of repeated STDP. Here again, latency decreases and stabilizes. Synapses are selected on the basis of their first-spike timings: the earliest ones are fully potentiated and the latest are weakened. The sum of synaptic weights follows the same behavior as in the spike wave experiment: it stabilizes around the output neuron's threshold value (see Figure 1B).

spikes. At this point, the system has converged: synaptic potential flattens around a value corresponding to the postsynaptic neuron's threshold (see Figure 2B, top).

Thus, the selection of the earliest spikes through STDP is obtained even when late parts of the spike trains carry the same amount of information.

3.2 Latency versus Firing Rates and Synchrony. Note that no assumptions were made in the previous simulation as to how neurons encode information within the spike trains. Firing rates (Gerstner, Kreiter, Markram, & Herz, 1997; Shadlen & Newsome, 1998) or synchrony (Abeles, 1991) are widely believed to support the neural code. In this regard, one may think that these principles should drive the effects of STDP. Indeed, it has been argued that competition for control of the postsynaptic response would thus be won by the most correlated inputs (Song, Miller, & Abbott, 2000), or by

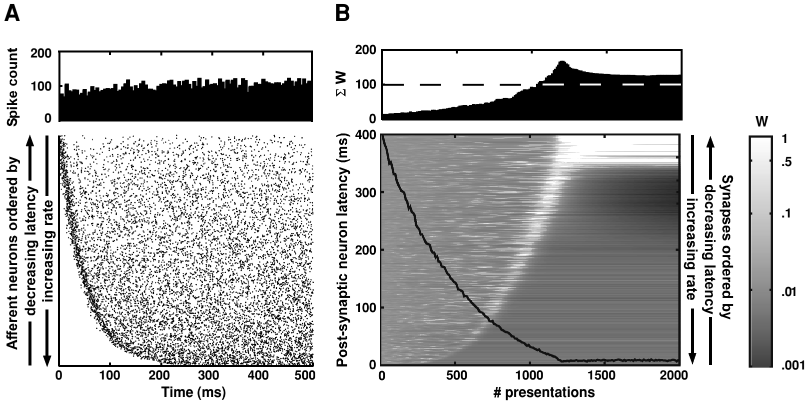

Figure 3: Latency versus rate. (A) Typical incoming activity. The shorter the latency, the fewer spikes the input neuron emits. According to the spike train design, the tail of the PSTH contains more spikes than its head. A jitter of 5 ms is applied to each spike timing, and no spontaneous activity is superimposed on the reproducible input structure so as to rigorously control the rates of afferents. Activity is then presented to the postsynaptic neuron, whose potential is set to 0. Prior to presentation, the presynaptic neurons do not fire any spike. (B) Dynamics of repeated STDP. The same trend emerges: although receiving fewer spikes, synapses hit by the earliest trains are finally fully potentiated. Inversely, strongly firing inputs will be neglected because they fire late.

the most vigorously firing ones (Gerstner & Kistler, 2002b). However, this study points toward first-spike timing as a determining factor, emphasizing the role of temporal asynchrony in neural information coding (Gautrais & Thorpe, 1998). In the brain, short latencies are generally associated with highest firing rates, which also often result in high temporal correlations. This would make it difficult to disentangle the respective influences of these aspects of the spike train on neuronal learning. In our simulations, however, we can ask how these different aspects fare, one compared to the other, by artificially defining input spike patterns where first-spike timing is pitted against average firing rate or amount of synchrony.

In the next simulation, latencies and rates have been artificially opposed: the afferent neuron with the shortest latency fires at the slowest rate over the entire window, the latest one at the highest rate, and neurons inbetween display a gradual latency-to-rate trade-off. As in the main experiment, spike times are jittered at each presentation (but spontaneous activity has been removed so as to rigorously control the firing rate of each input neuron, ranging from 4 to 44 Hz; (see Figure 3A). The result is clear as the same trend emerges (see Figure 3B). This means that a synapse receiving a very high firing rate will not be retained by STDP if it does not also correspond to one of the shortest latencies.

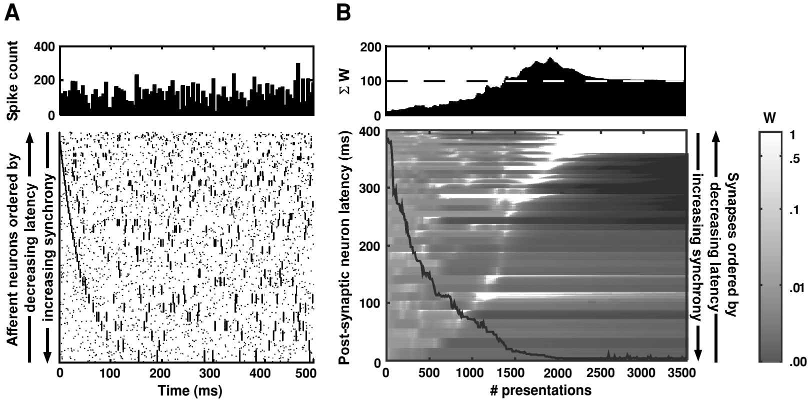

Figure 4: Latency versus synchrony. (A) Typical incoming activity. The spike pattern was designed so as to oppose latency to synchrony: the shorter the latency, the shorter the synfire chain. Spontaneous activity of 5 Hz has been superimposed on the reproducible input structure and no jitter is simulated so as to rigorously keep the synchrony within the spike trains. Activity is then presented to the postsynaptic neuron whose potential is set to 0. Prior to presentation, the presynaptic neurons do not fire any spike. (B) Dynamics of repeated STDP. Once again, postsynaptic latency decreases before stabilizing, and the earliest synapses become fully potentiated while later ones are weakened. Thus, being highly correlated is not sufficient for inputs to be selected by STDP; they also have to fire with one of the shortest latencies.

In the final simulation, latency and amount of synchrony are artificially opposed: the longer the latency, the more neurons are made to share the same spike train (hence, defining longer synfire chains; Abeles, 1991). The size of these correlated groups gradually extends from 1 to 54 synchronously firing units. Here, 5 Hz spontaneous activity is added as usual, while jitter is removed so as to obtain truly synchronous waves (see Figure 4A). Once again, the results are clear: synapses carrying the very first input spikes are selected by STDP, whereas highly correlated inputs having late latencies are depressed (see Figure 4B). Note the small and eventually vanishing streaks of synaptic potential, a necessary consequence of the fact that synchronous inputs are always reinforced or depressed together by the STDP rule.

Conclusively, when latency and synchrony are inversely correlated, the selection of the potentiated synapses still depends on the timing of the first reproducible spikes alone.

3.3 Selectivity Measures.

Convergence of synaptic weights onto earliest afferents through STDP is now established, underlining the importance of first-spikes timing. But why do synaptic weights distribute in such a remarkable way? The first obvious answer is the postsynaptic spike latency reduction, as displayed in all the simulations. But having a neuron responding fast to a given event is not an interesting feature if it is not also selective to it. STDP has been shown to explain the development of direction selectivity in recurrent cortical networks, where excitatory and inhibitory synapses were modified according to their spiking activity (Rao & Sejnowski, 2000). Selectivity measures were based on whether the postsynaptic neuron fires. How should one define selectivity within the present, inhibition-less framework? Any conventional measure based on the postsynaptic firing rate is ruled out since the postsynaptic response is limited to its first spike (see sections 2.1 and 4). Under these conditions, we propose to use firing latency as a viable measure of selectivity: if a neuron spikes faster to the input pattern it learned than to any other one, it would de facto behave if not selectively, at least differentially to it.

First, we generated 50,000 distractor spike trains using the same method (Poisson-inspired spike trains at 20 Hz; see section 3.1) as for the target pattern (see Figure 2A); a priori, these can be considered as equivalent to the target in terms of average latency of the spike trains, spike count, correlations, and so forth.

At each learning step of the main simulation (see Figure 2B), these distractor patterns were tested for postsynaptic latency on the learning neuron and compared to 1000 responses to the target pattern, using the same conditions of background noise (5 ms jitter and 5 Hz spontaneous activity). These distractors and target presentations did not trigger the learning rule: the postsynaptic latency was measured but did not give rise to any synaptic modifications. A threshold was set around the mean response time to targets. Target spike trains yielding responses before the threshold were considered as hits and distractor ones as false alarms. Selectivity (d') was computed as follows:

$$d' = z\,(\text{hit rate}) - z\,(\text{false alarm rate}),$$

where $z(p)$ stands for the inverse of the normal cumulative distribution function of p (Green & Swets, 1966). The expected maximum was computed using an expected hit rate of 50% and a false alarm rate of $1/2n$, where n is the number of distractor patterns (here, $n = 50,000$).

The results (see Figure 5) show that while the neuron was initially less likely to respond to the target pattern than to arbitrary distractors (due to its initially random weight distribution), it became highly selective to its target as learning developed: after about 1500 presentations, when the postsynaptic neuron latency and weights have indeed stabilized (see Figure 2B), not one of the distractor patterns could make this neuron fire sooner than with the target one.

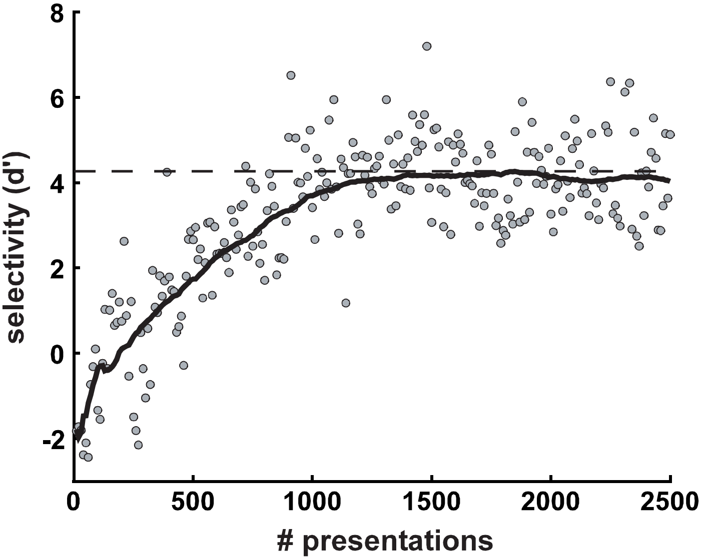

Figure 5: Selectivity of the efferent neuron. The dashed line is the expected maximum value for d' (under the present conditions), gray dots stand for d' values at each learning step, and the black curve is a local average of these. The neuron becomes steadily more selective to the target pattern used for training. Initially, it tends to fire slightly more slowly to it than to 50,000 spike trains used as distractors. After 1500 presentations, when its response latency stabilizes (see Figure 2B), the postsynaptic neuron reacts to its target before any of the distractors. Note that some observations landed above the expected maximum value for d': in certain conditions (100% hit rate and/or 0% false alarms), the observed d' can rise above the expected max value computed for an expected 50% hit rate and 1 for 100,000 false alarms.

Conclusively, a neuron exhibiting STDP will characteristically respond faster and faster to a precise repeated pattern than to any comparable one, thus becoming more selective to it, at least in the terms proposed here.

4 Discussion

The temporally asymmetric STDP rule reinforces the last synapses involved in making a neuron fire. At first sight, one might think that its repeated application invariably reinforces afferents at the tail of the spike sequence firing a neuron. Instead, we have demonstrated that it focuses, in a remarkably robust manner, on the head of the spike sequence: dynamical STDP tunes a neuron to synapses transmitting reproducible spikes with the shortest latencies, thus enabling the cell to respond faster and selectively to the input it has learned. This result is obtained for arbitrary structures of the input spike pattern.

This STDP tuning is achieved by selecting a relatively small subset of afferents, here fully determined by the model's output threshold. In that sense, it is worth noting that only 10 to 40 fully potentiated excitatory inputs, among a possible 10,000 or so, would be enough to evoke a response (Shadlen & Newsome, 1994).

The main assumption for the demonstration is that spike times within the input pattern must show a certain degree of reproducibility (which can be modulated by fairly high amounts of jitter and spontaneous activity). This is compatible with experimental observations in numerous brain structures (Heil, 1997; Petersen et al., 2001; Bair & Koch, 1996; Sestokas et al., 1991; Liu et al., 2001; Richmond & Optican, 1990; Victor & Purpura, 1996; Nakamura, 1998). This repetition of similar input patterns could quite simply be the result of multiple exposures with the same stimulus at different times in life. Alternatively, we propose that a single stimulus exposure could result in a sequence of similar processing waves through the rhythmic activity of cortical oscillations (Hopfield, 1995). Note that in both situations, one would not expect the state of the system to be the same from one step to another. In particular, the efferent neuron potential would be unlikely to have the same resting level every time the input pattern is presented. We have performed simulations where baseline fluctuations were simulated. As usual, the typical incoming activity was of the same kind as the one used in the main simulation: the reproducible input structure consisted of 1000 spike trains generated using a Poisson-inspired process (see section 3.1). At every step, a 5 ms jitter was applied to each precise spike time; then 5 Hz spontaneous activity was superimposed before being presented to the neuron. But instead of being systematically reset to 0, the postsynaptic neuron potential was set according to a gaussian process of mean 0 and width 50. Prior to presentation, the presynaptic neurons do not fire any spike. An identical convergence was reached, though less rapidly (see Figure 6).

For simplicity, all of the simulations so far have been conducted with a single postsynaptic spike, no matter what the activity is afterward: a single STDP learning rule was applied at each input pattern presentation. While it may seem a controversial simplification, experimental evidence suggests that the effect of the late spikes could be neglected (Froemke & Dan, 2002). As Tsodyks (2002) commented, "The main effect of STDP is well explained by the first pair of spikes, with the additional spike having only a marginal contribution." The first presynaptic spike is also the most relevant in evoking an excitatory postsynaptic potential (EPSP), because of synaptic depression (Thomson & Deuchars, 1994), a trend that is reinforced by synaptic redistribution, where a synapse depresses even more when potentiated (Markram & Tsodyks, 1996; see below for more details). Nonetheless, for completeness, we investigated the effects of having multiple postsynaptic spikes that repeatedly triggered the STDP mechanism. Here we implemented a potential

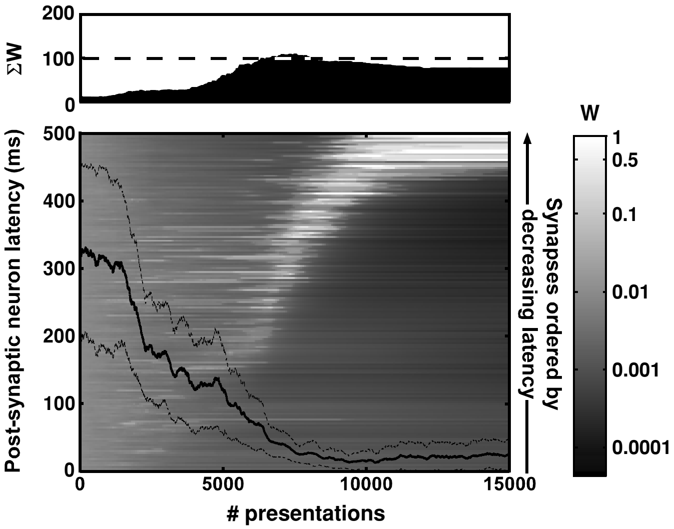

Figure 6: Fluctuating potential. In this simulation, the initial efferent neuron potential was set—just before the cell is presented with the incoming activity—according to a gaussian distribution of mean 0 and standard deviation 50. Typical incoming activity is of the same kind as in the main experiment: a reproducible structure of spike trains amid 5 ms jitter and 5 Hz spontaneous activity (see Figure 2A). Prior to presentation, the presynaptic neurons do not fire any spike. Due to the fluctuation, postsynaptic neuron latencies were highly variable; the black line thus depicts a smoothed version (using a moving average of width 101 steps; ± standard deviation, dotted lines). Here, the fluctuation results in a delay of the reinforcement process applied to the reproducible structure: from step 0 to 5000, the crest-ascending motion is not obvious. But once some synapses are strong enough (i.e., displayed in white on the figure), the modification sequence starts until only the earliest firing inputs remain with large weights and the postsynaptic neuron latency has decreased to reach a stable, steady state. The sum of synaptic weights follows the same behavior as in the spike wave experiment: it stabilizes around the output neuron's threshold value (see Figure 1B). While convergence needs more time to be reached (about 10 times more than for the original simulation; see Figure 2B), the repeated application of STDP leads the neuron with a highly fluctuating potential to tune itself on its earliest afferents and respond faster.

leak current ($\tau = 20$ ms) in the efferent neuron. Simulation parameters and typical incoming activity were of the same kind as in the main experiment (see Figure 2A). Again, the same trend was observed with a slowing of convergence and the appearance of irregular second postsynaptic spikes that did not prevent the neuron from becoming tuned to its earliest regular afferents (see Figure 7).

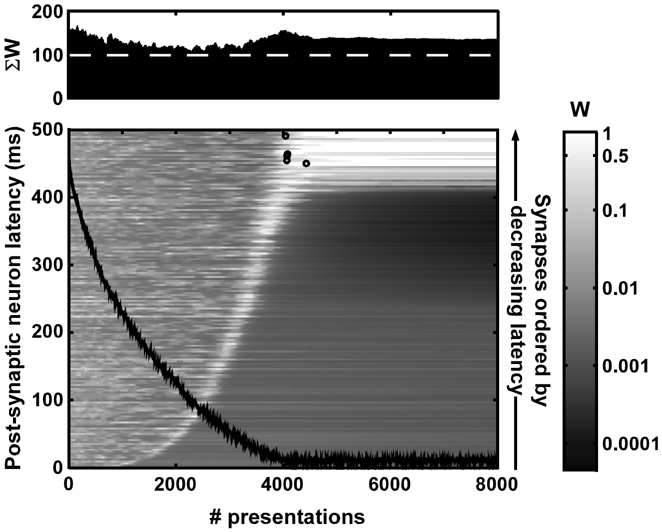

Figure 7: Leaky integrator and multiple postsynaptic responses. Here, a leak current ($\tau = 20$ ms) is added to the neuron potential and more than one postsynaptic spike can be elicited in the course of the presentation, triggering STDP repeatedly. Typical incoming activity is of the same kind as in the main experiment: a reproducible structure of spike trains amid 5 ms jitter and 5 Hz spontaneous activity (see Figure 2A). It is presented to the output neuron after setting its potential to 0. Prior to presentation, the presynaptic neurons do not fire any spike. The black line links the first postsynaptic response latency at each presentation. Times at which the efferent neuron has subsequent discharges are displayed by empty circles. Except for some additional postsynaptic responses that did not affect the dynamics of learning, convergence on the earliest afferents was reached, slightly more slowly than in the comparative case without the use of a leaky integrator (see Figure 2B, bottom). Due to the leakage current, synapses had to be initialized at a slightly higher value than in previous simulations. The sum of synaptic weights thus started from a higher value than before, above the threshold of the postsynaptic neuron's threshold (see Figure 2B, top). It nonetheless acted the same way to stabilize at threshold value.

One could also address the fact that theoretical studies use the idealized "smooth" STDP curve, while experimental data always display some noise: the curve is also likely to be noisy in biological settings (Bi & Poo, 1998). We thus tested this feature in a specific simulation under the same conditions as in section 3.1, except that each single synaptic modification was affected by a random offset taken from a gaussian distribution of mean 0 and with a standard deviation set to the noiseless amount of modification (see Figure 8A). Even when noise affects the synaptic modifications in a biologically realistic way, the trend emerges in a very similar manner (see Figure 8B).

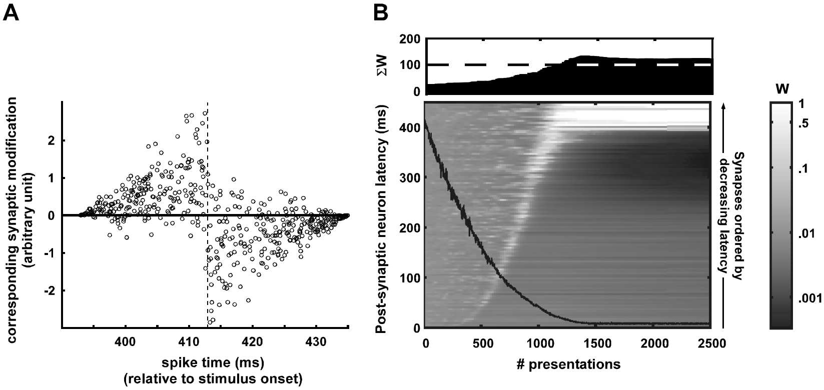

Figure 8: Noisy synaptic modifications. This simulation is identical to the main one (see Figure 2) except that in this case, each synaptic modification is affected by an offset taken from a gaussian distribution of mean 0 and standard deviation set to the noiseless amount of modification. Incoming activity is of the same kind as in the main experiment: a reproducible structure of spike trains amid 5 ms jitter and 5 Hz spontaneous activity (see Figure 2A). It is presented to the output neuron after setting its potential to 0. Prior to presentation, the presynaptic neurons do not fire any spike. (A) Typical sample of synaptic modifications. Taken from the modifications occurring at step 1 (only one out of two modifications are represented for clarity), this scatter plot shows the amount of each synaptic modification depending on the relative timing of the corresponding spike compared to the postsynaptic spike timing (dashed line). Notice how some expected rewards are in fact depressions and vice versa, as in the experimentally obtained STDP graphs (Bi & Poo, 1998). (B) Dynamics of repeated STDP. The trend emerges quite the same way as in Figure 2, illustrating the fact that STDP is able to reach for the first spikes of a reproducible input pattern, even when synaptic modifications are randomized.

Evidently, the demonstration here depends on the learning rule used in these simulations, and in particular on its temporally asymmetrical shape. This form was observed in several empirical studies (Markram et al., 1997; Bi & Poo, 1998, 2001; Zhang et al., 1998; Feldman, 2000; and Sjöström & Nelson, 2002, for reviews) and has triggered numerous theoretical studies of STDP (Kempter, Gerstner, & van Hemmen, 1999; Song, Miller, & Abbott, 2000; van Rossum, Bi, & Turrigiano, 2000; Gerstner & Kistler, 2002b; see Abbott & Nelson, 2000, and Kepecs, Van Rossum, Song & Tegner, 2002, for reviews). Other forms exist, for example, symmetric (Egger, Feldmeyer, & Sakmann, 1999) or opposed (Bell, Han, Sugawara, & Grant, 1997), that could give rise to radically different trends. In the latter case, for example, where a pre-before-post pairing would induce depression and post-before-pre reinforcement, inhibiting the last spikes to make a cell fire might in fact increase the postsynaptic latency from one presentation to the other.

A critical question concerns the handling of spike timing-dependent modifications when the same synapse receives more than one input spike that falls in the STDP window. Here we simply assumed that the effects of spike pairs would sum linearly, as in many theoretical studies (Gerstner, Kempter, van Hemmen, & Wagner, 1996; Song, Miller, & Abbott, 2000; Senn, Markram, & Tsodyks, 2001), but the question extends to the problem of synaptic modifications in natural spike trains. A recent experimental study showed that the efficacy of each spike in modifiying synaptic strength was suppressed by the preceding spike in the same neuron: STDP favors the earliest of two spike pairs (Froemke & Dan, 2002), that is, in our case, reinforcement over depression. Besides, in a spike sequence following a period of inactivity, each spike decreases the probability of vesicle release for the following presynaptic action potentials, reducing the amplitude of their respective EPSPs in the process (Thomson & Deuchars, 1994). This form of short-term plasticity, synaptic depression, acts at the presynaptic level through the depletion of the pool of vesicles at the synaptic release site, where recycling is a relatively slow process compared to sustained activity. Not only does it by itself suggest a relative importance of first spikes as opposed to later ones, but it also interacts with long-term potentiation in a supportive way called synaptic redistribution (Markram & Tsodyks, 1996). While increasing, or decreasing, the probability of transmitter release, STDP would at the same time decrease (resp. increase) the availability of readily releasable vesicles for later spikes: when they get stronger, synapses depress more and vice versa (Markram & Tsodyks, 1996; Volgushev, Voronin, Chistiakova, & Singer, 1997). As such, "synaptic redistribution can significantly enhance the amplitude of synaptic transmission for the first spikes in a sequence," thus giving a predominant role to the first spikes in terms of synaptic plasticity (Abbott & Nelson, 2000). Applied to the present framework, these more realistic simulations of synaptic modifications would no doubt reinforce the "back-in-time" dynamics of STDP, as illustrated by latency reduction, a necessary consequence of its very form.

5 Conclusion

The most important implications of our work concern the way the brain stores and uses information. Much work in neural coding assumes that information is encoded using either firing rate or synchrony. This study raises doubts about this assumption by demonstrating the unambigous prevalence of the earliest firing inputs. This sensitivity to spatiotemporal spike patterns, obviously inherent in STDP rules, emphasizes the importance of temporal structure in neural information, as proposed and argued by several authors (Perkel & Bullock, 1968; Bialek, Rieke, de Ruyter van Steveninck, & Warland, 1991; Engel, Konig, Kreiter, Schillen, & Singer, 1992; de Ruyter van Steveninck, Lewen, Strong, Koberle, & Bialek, 1997; Bair, 1999; Panzeri, Petersen, Schultz, Lebedev, & Diamond, 2001; Van Rullen & Thorpe, 2001). As stated by Froemke and Dan (2002), "Timing of the first spike in each burst is dominant in synaptic modifications." Latencies do indeed seem to play a distinct role in terms of neural processing, as shown throughout this study.

We showed how spatially distributed synapses can be made to integrate a spike wave coming from an afferent population and evoke a fast and selective response in the postsynaptic neuron, even in the presence of background noise. As a consequence, the fact that STDP naturally leads a neuron to respond rapidly and selectively on the basis of the first few spikes in its afferents lends support for the idea that even complex visual recognition tasks can be performed on the basis of a single wave of spikes (Van Rullen & Thorpe, 2002; Thorpe & Imbert, 1989).

Acknowledgments

This research was supported by the CNRS, the ACI Neurosciences Computationelles et Intégratives, and SpikeNet Technology SARL. We also thank our reviewers for their critical comments.

References

Abbott, L. F., & Nelson, S. B. (2000). Synaptic plasticity: Taming the beast. *Nat. Neurosci. 3*, (Suppl.), 1178–1183.

Abeles, M. (1991). *Corticonics: Neural circuits of the cerebral cortex*. Cambridge: Cambridge University Press.

Bair, W. (1999). Spike timing in the mammalian visual system. *Curr. Opin. Neurobiol., 9*, 447–453.

Bair, W., & Koch, C. (1996). Temporal precision of spike trains in extrastriate cortex of the behaving macaque monkey. *Neural Comput., 8*, 1185–1202.

Bell, C. C., Han, V. Z., Sugawara, Y., & Grant, K. (1997). Synaptic plasticity in a cerebellum-like structure depends on temporal order. *Nature, 387*, 278–281.

Bi, G. Q., & Poo, M. M. (1998). Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci., 18*, 10464–10472.

Bi, G. Q., & Poo, M. M. (2001). Synaptic modification by correlated activity: Hebb's postulate revisited. *Annu. Rev. Neurosci., 24*, 139–166.

Bialek, W., Rieke, F., de Ruyter van Steveninck, R. R., & Warland, D. (1991). Reading a neural code. *Science, 252*, 1854–1857.

Cateau, H., & Fukai, T. (2003). A stochastic method to predict the consequence of arbitrary forms of spike-timing-dependent plasticity. *Neural. Comput, 15*, 597–620.

de Ruyter van Steveninck, R. R., Lewen, G. D., Strong, S. P., Koberle, R., & Bialek, W. (1997). Reproducibility and variability in neural spike trains. *Science, 275*, 1805–1808.

Egger, V., Feldmeyer, D., & Sakmann, B. (1999). Coincidence detection and changes of synaptic efficacy in spiny stellate neurons in rat barrel cortex. *Nat. Neurosci., 2*, 1098–1105.

Engel, A. K., Konig, P., Kreiter, A. K., Schillen, T. B., & Singer, W. (1992). Temporal coding in the visual cortex: New vistas on integration in the nervous system. *TINS, 15*, 218-226.

Evarts, E. V. (1964). Temporal Patterns of Discharge of Pyramidal Tract Neurons During Sleep and Waking in the Monkey. *J. Neurophysiol., 27*, 152–171.

Feldman, D. E. (2000). Timing-based LTP and LTD at vertical inputs to layer II/III pyramidal cells in rat barrel cortex. *Neuron, 27*, 45–56.

Froemke, R. C., & Dan, Y. (2002). Spike-timing-dependent synaptic modification induced by natural spike trains. *Nature, 416*, 433–438.

Gautrais, J., & Thorpe, S. (1998). Rate coding versus temporal order coding: A theoretical approach. *Biosystems, 48*, 57–65.

Gerstner, W., Kempter, R., van Hemmen, J. L., & Wagner, H. (1996). A neuronal learning rule for sub-millisecond temporal coding. *Nature, 383*, 76–81.

Gerstner, W., & Kistler, W. M. (2002a). *Spiking neuron models.* Cambridge: Cambridge University Press.

Gerstner, W., & Kistler, W. M. (2002b). Mathematical formulations of Hebbian learning. *Biol. Cybern., 87*, 404–415.

Gerstner, W., Kreiter, A. K., Markram, H., & Herz, A. V. (1997). Neural codes: Firing rates and beyond. *Proc. Natl. Acad. Sci. USA, 94*, 12740–12741.

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics.* New-York: Wiley.

Heil, P. (1997). Auditory cortical onset responses revisited. I. First-spike timing. *J. Neurophysiol., 77*, 2616–2641.

Hopfield, J. J. (1995). Pattern recognition computation using action potential timing for stimulus representation. *Nature, 376*, 33–36.

Hubel, D. H. (1959). Single unit activity in striate cortex of unrestrained cats. *J. Physiol., 147*, 226–238.

Kempter, R., Gerstner, W., & van Hemmen, J. L. (1999). Hebbian learning and spiking neurons. *Phys. Rev. E, 59*, 4498–4514.

Kepecs, A., Van Rossum, M. C., Song, S., & Tegner, J. (2002). Spike-timing-dependent plasticity: Common themes and divergent vistas. *Biol. Cybern., 87*, 446–458.

Liu, R. C., Tzonev, S., Rebrik, S., & Miller, K. D. (2001). Variability and information in a neural code of the cat lateral geniculate nucleus. *J. Neurophysiol., 86*, 2789–2806.

Mainen, Z. F., & Sejnowski, T. J. (1995). Reliability of spike timing in neocortical neurons. *Science, 268*, 1503–1506.

Markram, H., Lubke, J., Frotscher, M., & Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science, 275*, 213–215.

Markram, H., & Tsodyks, M. (1996). Redistribution of synaptic efficacy between neocortical neurons. *Nature, 382*, 807–810.

Meister, M., & Berry, M. J., II. (1999). The neural code of the retina. *Neuron, 22* (3), 435–450.

Nakamura, K. (1998). Neural processing in the subsecond time range in the temporal cortex. *Neural Comput., 10*, 567–595.

Panzeri, S., Petersen, R. S., Schultz, S. R., Lebedev, M., & Diamond, M. E. (2001). The role of spike timing in the coding of stimulus location in rat somatosensory cortex. *Neuron, 29*, 769–777.

Perkel, D. H., & Bullock, P. H. (1968). Neural coding. *Neurosci. Res. Prog. Sum., 3*, 405–527.

Petersen, R. S., Panzeri, S., & Diamond, M. E. (2001). Population coding of stimulus location in rat somatosensory cortex. *Neuron, 32*, 503–514.

Rao, R. P. N., & Sejnowski, T. J. (2000). Predictive sequence learning in recurrent neocortical circuits. In S. A. Solla, T. K. Leen, & K.-R. Müller (Eds.), In *Advances in neural information processing systems, 12*. (pp 164–170).Cambridge, MA: MIT Press.

Richmond, B. J., & Optican, L. M. (1990). Temporal encoding of two-dimensional patterns by single units in primate primary visual cortex. II. Information transmission. *J. Neurophysiol., 64*, 370–380.

Rubin, J., Lee, D. D., & Sompolinsky, H. (2001). Equilibrium properties of temporally asymmetric Hebbian plasticity. *Phys. Rev. Lett., 86*, 364–367.

Senn, W., Markram, H., & Tsodyks, M. (2001). An algorithm for modifying neurotransmitter release probability based on pre- and postsynaptic spike timing. *Neural Comput., 13*, 35–67.

Sestokas, A. K., Lehmkuhle, S., & Kratz, K. E. (1991). Relationship between response latency and amplitude for ganglion and geniculate X- and Y-cells in the cat. *Int. J. Neurosci., 60*, 59–64.

Shadlen, M. N., & Newsome, W. T. (1994). Noise, neural codes and cortical organization. *Curr. Opin. Neurobiol., 4*, 569–579.

Shadlen, M. N., & Newsome, W. T. (1998). The variable discharge of cortical neurons: Implications for connectivity, computation, and information coding. *J. Neurosci., 18*, 3870–3896.

Sjostrom, P. J., & Nelson, S. B. (2002). Spike timing, calcium signals and synaptic plasticity. *Curr. Opin. Neurobiol., 12*, 305–314.

Song, S., Miller, K. D., & Abbott, L. F. (2000). Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nat. Neurosci., 3*, 919–926.

Steriade, M., Oakson, G., & Kitsikis, A. (1978). Firing rates and patterns of output and nonoutput cells in cortical areas 5 and 7 of cat during the sleep-waking cycle. *Exp. Neurol., 60*, 443–468.

Thomson, A. M., & Deuchars, J. (1994). Temporal and spatial properties of local circuits in neocortex. *Trends Neurosci., 17*, 119-126.

Thorpe, S. J., & Imbert, M. (1989). Biological constraints on connectionist models. In R. Pfeifer, Z. Schreter, & F. Fogelman-Souklié (Eds.), *Connectionism in perspective* (pp. 63–92). Amsterdam: Elsevier.

Tsodyks, M. (2002). Spike-timing-dependent synaptic plasticity—the long road towards understanding neuronal mechanisms of learning and memory. *Trends Neurosci., 25*, 599–600.

van Rossum, M. C., Bi, G. Q., & Turrigiano, G. G. (2000). Stable Hebbian learning from spike timing-dependent plasticity. *J. Neurosci., 20*, 8812–8821.

VanRullen, R., & Thorpe, S. J. (2001). Rate coding versus temporal order coding: What the retinal ganglion cells tell the visual cortex. *Neural Comput., 13*, 1255–1283.

VanRullen, R., & Thorpe, S. J. (2002). Surfing a spike wave down the ventral stream. *Vision Res., 42*, 2593–2615.

Victor, J. D., & Purpura, K. P. (1996). Nature and precision of temporal coding in visual cortex: A metric-space analysis. *J. Neurophysiol., 76*, 1310–1326.

Volgushev, M., Voronin, L. L., Chistiakova, M., & Singer, W. (1997). Relations between long-term synaptic modifications and paired-pulse interactions in the rat neocortex. *Eur. J. Neurosci., 9*, 1656–1665.

Zhang, L. I., Tao, H. W., Holt, C. E., Harris, W. A., and Poo, M. (1998). A critical window for cooperation and competition among developing retinotectal synapses. *Nature, 395* 37–44.

Received January 30, 2004; accepted August 25, 2004.