



Optimal Localist and Distributed Coding of Spatiotemporal Spike Patterns Through STDP and Coincidence Detection

Timothée Masquelier^{1,2*} and Saeed R. Kheradpisheh³

¹ Centre de Recherche Cerveau et Cognition, UMR5549 CNRS—Université Toulouse 3, Toulouse, France, ² Instituto de Microelectrónica de Sevilla (IMSE-CNM), CSIC, Universidad de Sevilla, Sevilla, Spain, ³ Department of Computer Science, Faculty of Mathematical Sciences and Computer, Kharazmi University, Tehran, Iran

OPEN ACCESS

Edited by:

Anthony N. Burkitt,
The University of Melbourne, Australia

Reviewed by:

Thomas Wennekers,
Plymouth University, United Kingdom
Florence Isabelle Kleberg,
Frankfurt Institute for Advanced
Studies, Germany

*Correspondence:

Timothée Masquelier
timothee.masquelier@cnrs.fr

Received: 26 June 2018

Accepted: 17 August 2018

Published: 18 September 2018

Citation:

Masquelier T and Kheradpisheh SR
(2018) Optimal Localist and
Distributed Coding of Spatiotemporal
Spike Patterns Through STDP and
Coincidence Detection.
Front. Comput. Neurosci. 12:74.
doi: 10.3389/fncom.2018.00074

Repeating spatiotemporal spike patterns exist and carry information. Here we investigated how a single spiking neuron can optimally respond to one given pattern (localist coding), or to either one of several patterns (distributed coding, i.e., the neuron's response is ambiguous but the identity of the pattern could be inferred from the response of multiple neurons), but not to random inputs. To do so, we extended a theory developed in a previous paper (Masquelier, 2017), which was limited to localist coding. More specifically, we computed analytically the signal-to-noise ratio (SNR) of a multi-pattern-detector neuron, using a threshold-free leaky integrate-and-fire (LIF) neuron model with non-plastic unitary synapses and homogeneous Poisson inputs. Surprisingly, when increasing the number of patterns, the SNR decreases slowly, and remains acceptable for several tens of independent patterns. In addition, we investigated whether spike-timing-dependent plasticity (STDP) could enable a neuron to reach the theoretical optimal SNR. To this aim, we simulated a LIF equipped with STDP, and repeatedly exposed it to multiple input spike patterns, embedded in equally dense Poisson spike trains. The LIF progressively became selective to every repeating pattern with no supervision, and stopped discharging during the Poisson spike trains. Furthermore, tuning certain STDP parameters, the resulting pattern detectors were optimal. Tens of independent patterns could be learned by a single neuron using a low adaptive threshold, in contrast with previous studies, in which higher thresholds led to localist coding only. Taken together these results suggest that coincidence detection and STDP are powerful mechanisms, fully compatible with distributed coding. Yet we acknowledge that our theory is limited to single neurons, and thus also applies to feed-forward networks, but not to recurrent ones.

Keywords: neural coding, localist coding, distributed coding, coincidence detection, leaky integrate-and-fire neuron, spatiotemporal spike pattern, unsupervised learning, spike-timing-dependent plasticity (STDP)

1. INTRODUCTION

In a neural network, either biological or artificial, two forms of coding can be used: localist or distributed. With localist coding, each neuron codes (i.e., maximally responds) for one and only one category of stimulus (or stimulus feature). As a result, the category of the stimulus (or the presence of a certain feature) can be inferred from the response of this sole neuron, ignoring the other neurons' responses. Conversely, with distributed coding each neuron responds to multiple stimulus categories (or features) in a similar way. Therefore, the response of each neuron is ambiguous, and the category of the stimulus, or the presence of a certain feature, can only be inferred from the responses of multiple neurons. Thus the distinction between the two schemes is the number of different stimuli to which a given neuron responds—not the number of neurons which respond to a given stimulus. Indeed, a localist network can have redundancy, and use multiple “copies” of each category specific neuron (Thorpe, 1989; Bowers, 2009).

Does the brain use localist or distributed coding? This question has been, and still is, intensively debated. In practice, discriminating between the two schemes from electrophysiological recordings is tricky (Quiroga and Kreiman, 2010), since the set of tested stimuli is always limited, the responses are noisy, the thresholds are arbitrary and the boundaries between categories are fuzzy. Here we do not attempt to do a complete review of the experimental literature; but rather to summarize it. It is commonly believed that distributed coding is prevalent (Rolls et al., 1997; O'Reilly, 1998; Hung et al., 2005; Quiroga et al., 2008), but there is also evidence for localist coding, at least for familiar stimuli, reviewed in Bowers (2009, 2017), Thorpe (2009, 2011), and Roy (2017).

The question of localist vs. distributed coding is also relevant for artificial neural networks, and in particular for the recently popular deep neural networks. Most of the time, these networks are trained in a supervised manner, using the backpropagation algorithm (LeCun et al., 2015). The last layer contains exactly one neuron per category, and backpropagation forces each neuron to respond more strongly when the stimulus belongs to the neuron's category. In other words, localist coding is imposed in the last layer. Conversely, the hidden layers are free to choose their coding scheme, which is supposedly optimal for the categorization task at hand. It is thus very interesting to analyze the chosen coding scheme. It is not easy to do such analysis on the brain (as explained above), but we can do it rigorously for computational models by computing the responses to huge amounts of images, and even synthesizing images that maximize the responses. Results indicate that some hidden neurons respond to one object category only (Zhou et al., 2015; Nguyen et al., 2016; Olah et al., 2017), while others respond to multiple different objects (Nguyen et al., 2016; Olah et al., 2017). Thus it appears that both localist and distributed codes can be optimal, depending on the task, the layer number, and the network parameters (number of layers, neurons, etc.).

Let us come back to the brain, in which computation is presumably implemented by spiking neurons performing coincidence detection (Abeles, 1982; König et al., 1996; Brette,

2015). This observation raises an important question, which we tried to address in this theoretical paper: can coincidence detector neurons implement both localist and distributed codes? In this context, different stimuli correspond to different spatiotemporal input spike patterns. Here each pattern was generated randomly, leading to chance-level overlap between patterns. In addition, each pattern was jittered at each presentation, resulting in categories of similar, yet different, stimuli. Can a neuron respond to one, or several of these patterns, and not to random inputs? What is the required connectivity to do so in an optimal way? And finally, can this required connectivity emerge with spike-timing-dependent plasticity (STDP), in an unsupervised manner?

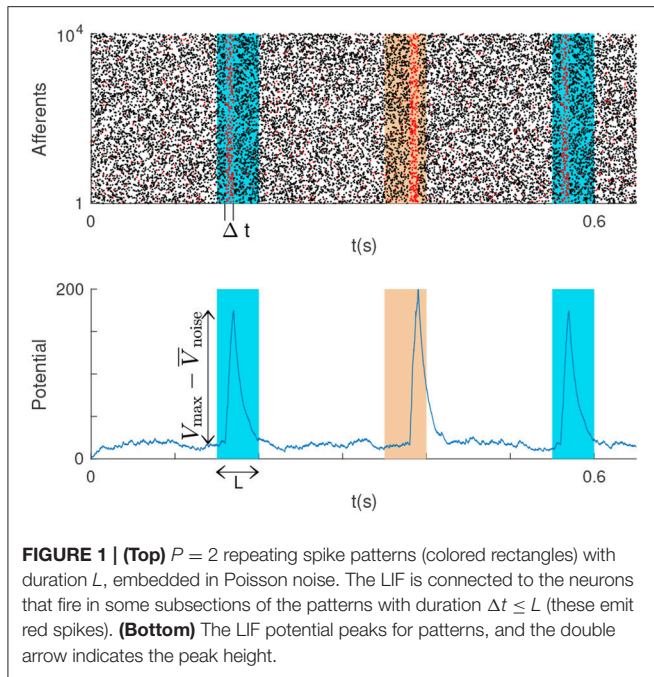
To address these questions, we extended a theory that we developed in a previous paper, but which was limited to one pattern only, i.e., localist coding (Masquelier, 2017), to the multi-pattern case. Briefly, we derived analytically the signal-to-noise ratio (SNR) of a multi-pattern detector, and investigated the conditions for its optimality. In addition, using numerical simulations, we showed that a single neuron equipped with STDP can become selective to multiple repeating spike patterns, even without supervision and that the resulting detectors can be close to the theoretical optimum. Surprisingly, a single neuron could robustly learn up to ~ 40 independent patterns (using parameters arguably in the biological range). This was not clear from previous simulations studies, in which neurons equipped with STDP only learned one pattern (localist coding) (Masquelier et al., 2008, 2009; Gilson et al., 2011; Humble et al., 2012; Hunzinger et al., 2012; Kasabov et al., 2013; Klampfl and Maass, 2013; Nessler et al., 2013; Krunglevicius, 2015; Sun et al., 2016; Masquelier, 2017), or two patterns (Yger et al., 2015). This shows that STDP and coincidence detection are compatible with distributed coding.

2. FORMAL DESCRIPTION OF THE PROBLEM

The problem we addressed is similar to the one of Masquelier (2017), but extended to the multi-pattern case. For the reader's convenience, we fully describe it below.

We addressed the problem of detecting one or several spatiotemporal spike patterns with a single LIF neuron. Intuitively, one should connect the neurons that are active during the patterns (or during subsections of them) to the LIF neuron. That way, the LIF will tend to be more activated by the patterns than by some other inputs. More formally, we note P the number of spike patterns, and assume that they all have the same duration L . We note N the number of neurons involved. For each pattern, we chose a subsection with duration $\Delta t \leq L$, and we connect the LIF to the M neurons that emit at least one spike during at least one of these subsections (Figure 1).

We hypothesize that all afferent neurons fire according to a homogeneous Poisson process with rate f , both inside and outside the patterns. That is the patterns correspond to some realizations of the Poisson process, which can be repeated (this is sometimes referred to a “frozen noise”). At each repetition



a random time lag (jitter) is added to each spike, drawn from a uniform distribution over $[-T, T]$ (a normal distribution is more often used, but it would not allow analytical treatment, Masquelier, 2017).

We also assume that synapses are instantaneous, which facilitates the analytic calculations.

For now we ignore the LIF threshold, and we want to optimize its signal-to-noise ratio (SNR), defined as:

$$SNR = \frac{V_{\max} - \bar{V}_{\text{noise}}}{\sigma_{\text{noise}}}, \quad (1)$$

where V_{\max} is the maximal potential reached during the pattern presentations, \bar{V}_{noise} is the mean value for the potential with Poisson input (noise period), and σ_{noise} is its standard deviation. Obviously, a higher SNR means a larger difference between the LIF membrane potential during the noise periods and its maximum value, which occurs during the selected Δt window of each pattern. Therefore, the higher the SNR the lower the probability of missing patterns, and of false alarms.

We consider that P, L, N, f , and T are imposed variables, and that we have the freedom to choose $\Delta t \leq L$ and the membrane time constant τ in order to maximize the SNR.

We note that this problem is related to the synfire chain theory (Abeles, 1991). A synfire chain consists of a series of pools of neurons linked together in a feed-forward chain, so that volleys of synchronous spikes can propagate from pool to pool in the chain. Each neuron can participate in several of such chains. The number of different chains that can coexist in a network of a given size has been termed capacity. This capacity can be optimized (Herrmann et al., 1995). To do so, a given neuron should respond to certain spike volleys, but not to others, which is similar to our optimization of a multi-pattern SNR. Yet it is

also different: we use homogeneous Poisson activity, not spike volleys, and we ignore the threshold, while synfire chains require thresholds.

3. A THEORETICAL OPTIMUM

3.1. Deriving the SNR Analytically

Here we are to find the optimum SNR of the LIF for P patterns. To this end we should first calculate the SNR analytically. Again, the derivations are similar to the ones in Masquelier (2017), but extended to the multi-pattern case (which turned to mainly impact Equation 7).

In this section, we assume non-plastic unitary synaptic weights. That is an afferent can be either connected ($w = 1$) or disconnected ($w = 0$) [in the **Appendix** (Supplementary Material) we estimate the cost of this constraint on the SNR]. Thus the LIF obeys the following differential equation:

$$\tau \frac{dV}{dt} = -V + \tau \sum_i \delta(t - t_i), \quad (2)$$

where the t_i are the presynaptic spike times of all the connected afferents.

Since synapses are instantaneous and firing is Poissonian, during the noise periods and outside the Δt windows we have: $\bar{V}_{\text{noise}} = \tau f M$ and $\sigma_{\text{noise}} = \sqrt{\tau f M / 2}$ (Burkitt, 2006), where M is the number of connected input neurons (with unitary weights).

To compute V_{\max} , it is convenient to introduce the reduced variable:

$$v_{\max} = \frac{V_{\max} - \bar{V}_{\text{noise}}}{\bar{V}^{\infty} - \bar{V}_{\text{noise}}}, \quad (3)$$

where $\bar{V}^{\infty} = \tau r$ is the mean potential of the steady regime that would be reached if Δt was infinite, and r is the input spike rate during the Δt window, resulting from the total received spikes from all input neurons during this window.

v_{\max} can be calculated by exact integration of the LIF differential equation (Masquelier, 2017). Here we omit the derivation and present the final equation:

$$v_{\max} = \min \left(1, \frac{\Delta t}{2T} \right) - \frac{\tau}{2T} \log \left(1 - e^{-\max(\Delta t, 2T)/\tau} + e^{-|\Delta t - 2T|/\tau} \right). \quad (4)$$

Using the definition of v_{\max} in Equation (3), we can rewrite the SNR equation as:

$$SNR = v_{\max} \frac{\bar{V}^{\infty} - \bar{V}_{\text{noise}}}{\sigma_{\text{noise}}}. \quad (5)$$

Obviously, different Poisson pattern realizations will lead to different values for M and r that consequently affect each of the terms \bar{V}^{∞} , \bar{V}_{noise} , and σ_{noise} . Here we want to compute

the expected value of the SNR across different Poisson pattern realizations:

$$\begin{aligned} \langle SNR \rangle &= v_{\max} \left\langle \frac{\overline{V}^{\infty} - \overline{V}_{\text{noise}}}{\sigma_{\text{noise}}} \right\rangle \\ &= v_{\max} \sqrt{2\tau/f} \left\langle \frac{r - f\langle M \rangle}{\sqrt{\langle M \rangle}} \right\rangle \\ &\approx v_{\max} \sqrt{2\tau/f} \frac{\langle r \rangle - f\langle M \rangle}{\sqrt{\langle M \rangle}}. \end{aligned} \quad (6)$$

In section 3.2 we justify this last approximation through numerical simulations, and we also show that this average SNR is not much different from the SNR of particular Poisson realizations.

The last step to compute $\langle SNR \rangle$ in Equation 6 is to calculate $\langle M \rangle$ and $\langle r \rangle$. Since firing is Poissonian with rate $\lambda = f\Delta t$, the probability that a given afferent fires at least once in a given pattern subsection of length Δt is $p = 1 - e^{-f\Delta t}$. Here, we consider independent patterns, i.e., with chance-level overlap. Hence the probability that a given afferent fires at least once in at least one of the P pattern subsection is $1 - (1 - p)^P$. Thus the number of selected afferents M is on average:

$$\langle M \rangle = N(1 - (1 - p)^P) = N(1 - e^{-Pf\Delta t}). \quad (7)$$

Finally, the expected effective input spike rate during the Δt window is the expected total number of spikes, $fN\Delta t$, divided by Δt , thus:

$$\langle r \rangle = fN. \quad (8)$$

We note that the SNR scales with \sqrt{N} . In the rest of this paper we used $N = 10^4$ afferents, which is in the biological range.

3.2. Numerical Validations

We first checked if the variability of the SNR across Poisson realizations is small, and also if the approximation we made to compute the average SNR in Equation 6 is reasonable. To this aim, we generated 10^5 Poisson patterns, and computed M , r and the reduced SNR, $snr = (\langle r \rangle - f\langle M \rangle)/\sqrt{\langle M \rangle}$, for each of them (i.e., the right factor of the SNR in Equation 6, which is the only one that depends on the Poisson realization). As can be seen on **Figure 2**, left, M and r are strongly correlated, and the data points lie near a line which corresponds to nearly constant snr values (see the colored background). In other words, the snr does not change much for different Poisson pattern realizations and the average snr well represents the snr distribution even for the worst and best cases.

In addition, as can be seen on **Figure 2**, right, the average snr across different Poisson patterns is very close to the snr corresponding to the average-case scenario, i.e., $M = \langle M \rangle$ and $r = \langle r \rangle$ (as defined by Equations 7 and 8, respectively). Note that this Figure was done with relatively small values for the parameters P , Δt and f (respectively 1, 2 ms, and 1 Hz). Our simulations indicate that when increasing these parameter values, the approximation becomes even better (data not shown).

Next, we verified the complete SNR formula (Equation 6), which also includes v_{\max} , through numerical simulations. We

used a clock-based approach, and integrated the LIF equation using the forward Euler method with a 0.1 ms time bin. We used $P = 1$ and $P = 5$ patterns, and performed 100 simulations with different random Poisson patterns of duration $L = 20$ ms with rate $f = 5$ Hz. We chose $\Delta t = L = 20$ ms, i.e., the LIF was connected to all the afferents that emitted at least once during one of the patterns. In order to estimate V_{\max} , each pattern was presented 1,000 times, every 400 ms. The maximal jitter was $T = 5$ ms. Between pattern presentations, the afferents fired according to a Poisson process, still with rate $f = 5$ Hz, which allowed to estimate $\overline{V}_{\text{noise}}$ and σ_{noise} . We could thus compute the SNR from Equation 1 (and its standard deviation across the 100 simulations), which, as can be seen on **Figure 3**, matches very well the theoretical values, for $P = 1$ and 5. Note that the SNR standard deviation is small, which confirms that the average SNR, i.e., $\langle SNR \rangle$, represents well the individual ones.

3.3. Optimizing the SNR

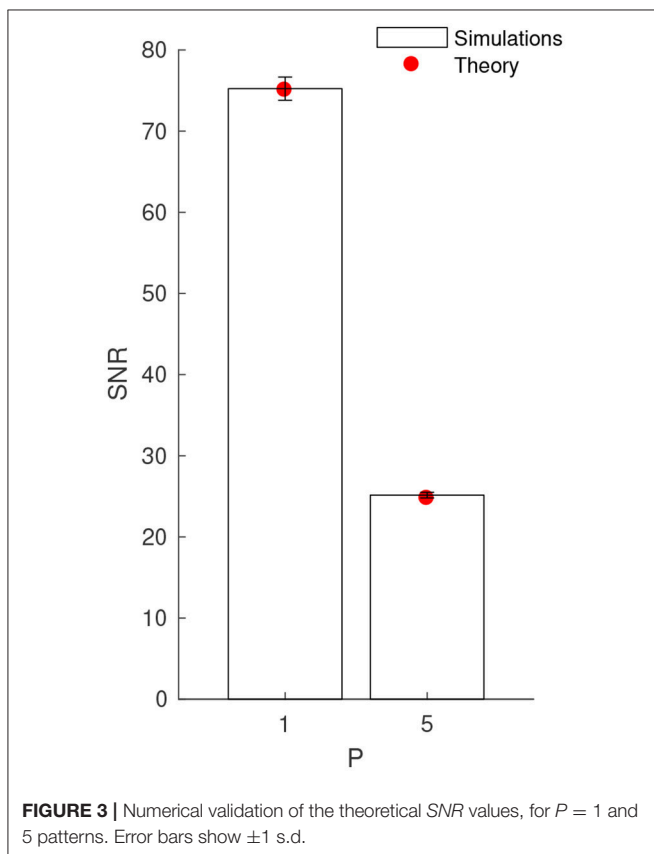
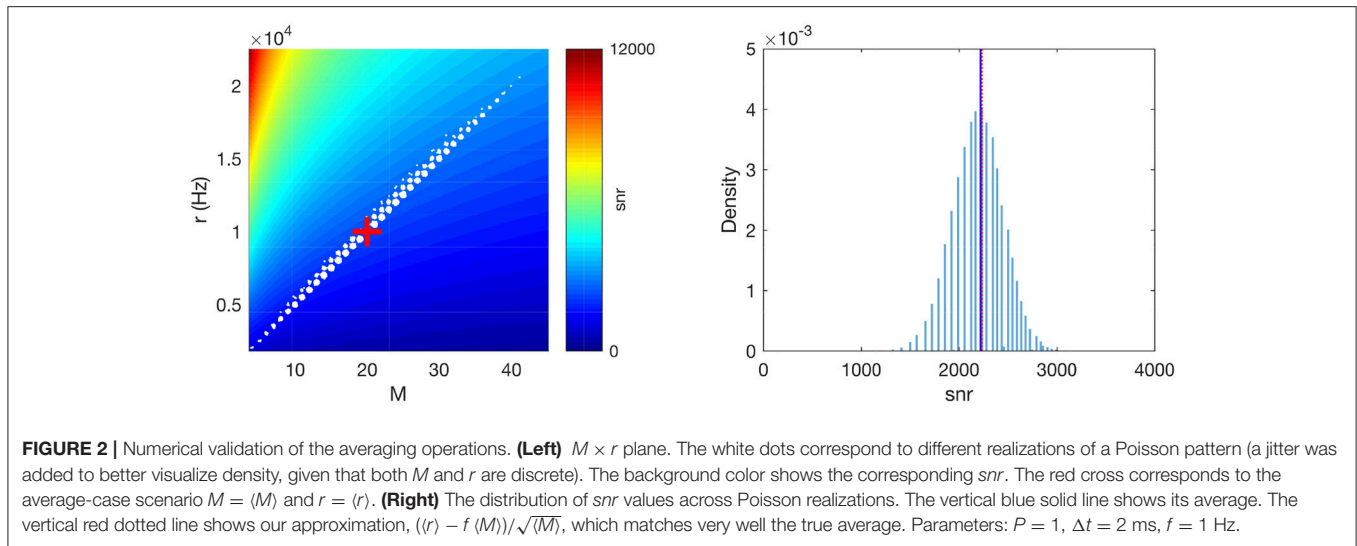
We now want to optimize the SNR given by Equation 6, by tuning τ and Δt . We also add the constraint $\tau f M \geq 10$ (large number of synaptic inputs), so that the distribution of V is approximately Gaussian (Burkitt, 2006). Otherwise, it would be positively skewed¹, thus a high SNR would not guarantee a low false alarm rate. We assume that L is sufficiently large so that an upper bound for Δt is not needed. We used the Matlab R2017a Optimization Toolbox (MathWorks Inc., Natick, MA, USA) to compute the optimum numerically.

Figure 4 illustrates the results with $P = 2$. One can make the following observations (similar to our previous paper which was limited to $P = 1$; Masquelier, 2017):

- Unless f and T are both high, the optimal τ and Δt have the same order of magnitude (see **Figure 4**, left).
- Unless T is high (>10 ms), or f is low (<1 Hz), then these timescales should be relatively small (at most a few tens of ms; see **Figure 4**, middle). This means that even long patterns (hundreds of ms or more) are optimally detected by a coincidence detector working at a shorter timescale, and which thus ignores most of the patterns. One could have thought that using $\tau \sim L$, to integrate all the spikes from the pattern would be the best strategy. But a long τ also decreases the detector's temporal resolution, thus patterns and random inputs elicit more similar responses, decreasing the SNR. Hence there is a trade-off, and it turns out that it is often more optimal to have $\tau < L$, that is to use subpatterns as signatures for the whole patterns.
- Unsurprisingly, the optimal SNR decreases with T (see **Figure 4**, right). What is less trivial, is that it also decreases with f . In other words, sparse activity is preferable. We will come back to this point in the discussion.

What is the biological range for T , which corresponds to the spike time precision? Millisecond precision in cortex has been reported (Kayser et al., 2010; Panzeri and Diamond, 2010; Havenith et al., 2011). We are aware that other studies found

¹With a low number of synaptic inputs, the mean V is close to zero. Since V is non-negative, its distribution is not symmetric anymore, but positively skewed.



poorer precision, but this could be due to uncontrolled variable or the use of inappropriate reference times (Masquelier, 2013).

In the rest of the paper we focus, as an example, on the point on the middle of the $T \times f$ plane $-T = 3.2$ ms and $f = 3.2$ Hz. When increasing P , the optimal τ and Δt decrease (Figure 5). Unsurprisingly, the resulting SNR also decreases,

but only slowly. It thus remains acceptable for several tens of independent patterns (e.g., $SNR \sim 7$ for $P = 40$).

4. SIMULATIONS SHOW THAT STDP CAN BE CLOSE-TO-OPTIMAL

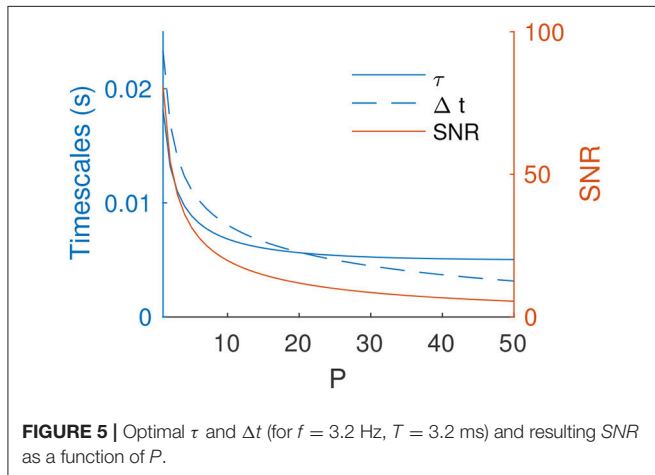
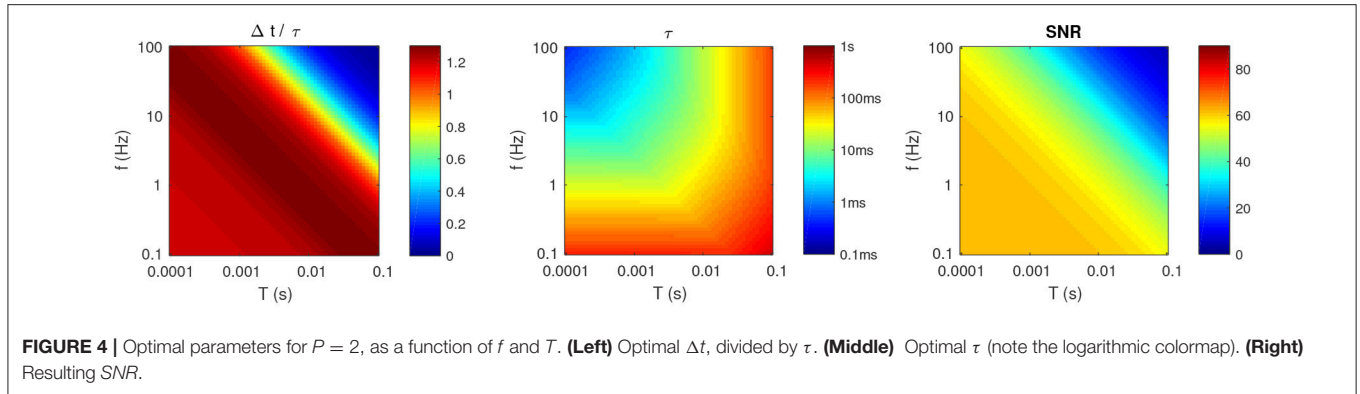
Next we investigated, through numerical simulations, if STDP could turn a LIF neuron into an optimal multi-pattern detector. More specifically, since STDP does not adjust the membrane time constant τ , we set it to the optimal value and investigated whether STDP could learn all the patterns with an optimal Δt^2 . Here, unlike in the previous section, we had to introduce a threshold, in order to have postsynaptic spikes, which are required for STDP. As a result, the optimal V_{max} , computed in the previous section, was never reached. Yet a high V_{max} guarantees a low miss rate, and a low \bar{V}_{noise} guarantees a low false alarm rate. Optimizing the previously defined SNR thus makes sense.

Again, we used a clock-based approach, and the forward Euler method with a 0.1 ms time bin. The Matlab R2017a code for these simulations has been made available in ModelDB (Hines et al., 2004) at <https://senselab.med.yale.edu/modeldb/>.

4.1. Input Spikes

The setup we used was similar to the one of our previous studies (Masquelier et al., 2008, 2009; Gilson et al., 2011; Masquelier, 2017). Between pattern presentations, the input spikes were generated randomly with a homogeneous Poisson process with rate f . The P spike patterns with duration $L = 100$ ms were generated only once using the same Poisson process (frozen noise). The pattern presentations occurred every 400 ms [in previous studies, we demonstrated that irregular intervals did

²When L is large (say tens of ms), STDP will typically not select all the afferents that fire in a full pattern, but only those that fire in a subsection of it, typically located at the beginning (Masquelier et al., 2008; Gilson et al., 2011; Masquelier, 2017), unless competition forces the neurons to learn subsequent subsections (Masquelier et al., 2009). The subsection duration depends on the parameters, and here we investigate the conditions under which this duration is optimal.



not matter (Masquelier et al., 2008, 2009; Gilson et al., 2011), so here regular intervals were used for simplicity]. The P patterns were presented alternatively, over and over again. **Figure 6** shows an example with $P = 2$ patterns. At each pattern presentation, all the spike times were shifted independently by some random jitters uniformly distributed over $[-T, T]$.

4.2. A LIF Neuron With Adaptive Threshold

We simulated a LIF neuron connected to all of the N afferents with plastic synaptic weights $w_i \in [0, 1]$, thus obeying the following differential equation:

$$\tau \frac{dV}{dt} = -V + \tau \sum_{ij} w_i(t_{ij}) \delta(t - t_{ij}), \quad (9)$$

where t_{ij} is the time of the j^{th} spike of afferent i .

We used an adaptive threshold [unlike in our previous studies (Masquelier et al., 2008, 2009; Gilson et al., 2011; Masquelier, 2017), in which a fixed threshold was used]. This adaptive threshold was increased by a fixed amount ($1.8\theta_0$) at each postsynaptic spike, and then exponentially decayed toward its baseline value θ_0 with a time constant $\tau_\theta = 80$ ms. This is a simple, yet good model of cortical cells, in the sense that it predicts very well the spikes elicited by a given input

current (Gerstner and Naud, 2009; Kobayashi et al., 2009). Here, such an adaptive threshold is crucial to encourage the neuron to learn multiple patterns, as opposed to fire multiple successive spikes to the same pattern. Since the theory developed in the previous sections ignored the LIF threshold, using an adaptive one is not worse than a fixed one, in the sense that it does not make the theory less valid.

We did not know which value for θ_0 could lead to the optimum. We thus performed an exhaustive search, using a geometric progression with a ratio of 2.5%.

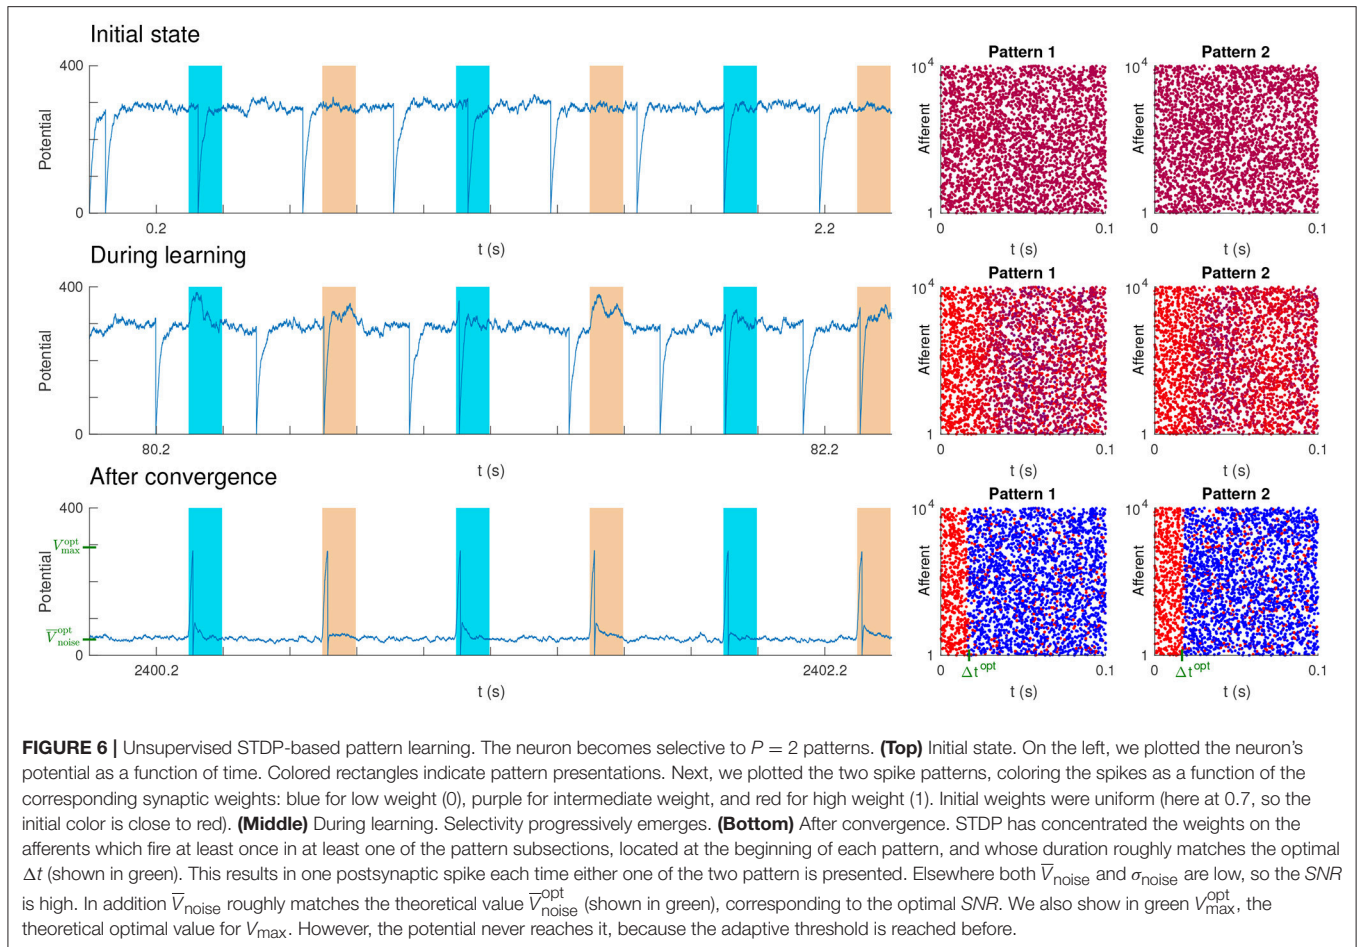
4.3. Synaptic Plasticity

Initial synaptic weights were all equal. Their value was computed so that $\bar{V}_{noise} = \theta + \sigma_{noise}$ (leading to an initial firing rate of about 4 Hz, see **Figure 6**, top). They then evolved in $[0, 1]$ with all-to-all spike STDP. Yet, we only modeled the Long Term Potentiation (LTP) part of STDP, ignoring its Long Term Depression (LTD) term. As in Song et al. (2000), we used a trace of presynaptic spikes at each synapse i , A_{pre}^i , which was incremented by δA_{pre} at each presynaptic spike, and then exponentially decayed toward 0 with a time constant $\tau_{pre} = 20$ ms. At each postsynaptic spike this trace was used for LTP at each synapse: $w_i \rightarrow w_i + w_i(1 - w_i)A_{pre}^i$.

Here LTD was modeled by a simple homeostatic mechanism. At each postsynaptic spike, all synapses were depressed: $w_i \rightarrow w_i + w_i(1 - w_i)w^{out}$ where $w^{out} < 0$ is a fixed parameter (Kempster et al., 1999).

Note that for both LTP and LTD we used the multiplicative term $w_i(1 - w_i)$, in contrast with additive STDP, with which the Δw is independent of the current weight value (Kempster et al., 1999; Song et al., 2000). This multiplicative term ensures that the weights remain in the range $[0, 1]$, and the weight dependence creates a soft bound effect: when a weight approaches a bound, weight changes tend toward zero. Here it was found to increase performance (convergence time and stability), in line with our previous studies (Masquelier and Thorpe, 2007; Kheradpisheh et al., 2016, 2018; Mozafari et al., 2018b).

The ratio between LTP and LTD, that is between δA_{pre} and w^{out} is crucial: the higher, the more synapses are maximally potentiated ($w = 1$) after convergence. Here we chose to keep $\delta A_{pre} = 0.1$ and to systematically vary w^{out} , using again a geometric progression with a ratio of 2.5%.



4.4. Results

For each $\theta_0 \times w^{\text{out}}$ point, 100 simulations were performed with different random pattern realizations, and we computed the proportion of “optimal” ones (see below), and reported it in **Table 1**. After 12,000 s of simulated time, the synaptic weights had all converged by saturation. That is synapses were either completely depressed ($w = 0$), or maximally potentiated ($w = 1$). A simulation was considered optimal if

1. all the patterns were learned, and
2. in an optimal way, that is if all patterns exhibited a subsection in which all spikes corresponded to maximally potentiated synapses ($w = 1$), and whose duration roughly matched the theoretical optimal Δt . In practice, we used the total number of potentiated synapses as a proxy of the mean subsection duration (since there is a non-ambiguous mapping between the two variables, given by Equation 7), and checked if this number matched the theoretical optimal M (Equation 7) with a 5% margin.

Note that this second condition alone would be easy to satisfy: the total amount of potentiated synapses is determined by the LTP/LTD ratio which we adjusted by fine-tuning w^{out} . However, satisfying the two conditions is harder, especially when P increases (**Table 1**).

It is worth mentioning that the learned subsections always corresponded to the beginning of the patterns, because STDP tracks back through them (Masquelier et al., 2008, 2009; Gilson et al., 2011), but this is irrelevant here since all the subsections are equivalent for the theory. **Figure 6** shows an optimal simulation with $P = 2$ patterns.

As can be seen in **Table 1**, the proportion of optimal simulations decreases with P , as expected. But more surprisingly, several tens of patterns can be optimally learned with reasonably high probability. With $P = 40$ the probability of optimal simulations is only 58%, but the average number of learned patterns is high: 39.5! This means that nearly all patterns are learned in all simulations, yet sometimes in a suboptimal manner. Finally, **Figure 7** shows that convergence time increases with P .

5. DISCUSSION

The fact that STDP can generate selectivity to any repeating spike pattern in an unsupervised manner is a remarkable, yet well documented fact (Masquelier et al., 2008, 2009; Gilson et al., 2011; Humble et al., 2012; Hunzinger et al., 2012; Kasabov et al., 2013; Klampfl and Maass, 2013; Nessler et al., 2013; Krunglevicius, 2015; Yger et al., 2015; Sun et al., 2016; Masquelier, 2017). Here we have shown that, surprisingly, a single neuron can become

TABLE 1 | Performance as a function of the number of patterns P .

P	5	10	20	40
Δt^{opt} (ms)	11	8.1	5.7	3.7
τ^{opt} (ms)	8.9	6.8	5.6	5.1
M^{opt}	1,600	2,300	3,100	3,800
SNR^{opt}	31	20	12	6.7
θ_0	190	140	110	92
w^{out}	$-6.2 \cdot 10^{-3}$	$-6.3 \cdot 10^{-3}$	$-6.5 \cdot 10^{-3}$	$-6.7 \cdot 10^{-3}$
$\langle P_{\text{learned}} \rangle$	5	10	20	39.5
Hit rate (%)	98.9	98.6	97.9	96.5
False alarms (Hz)	0	0	0	0
$P(\text{opt})$ (%)	100	100	100	58

The first four lines are computed from the theoretical optimum. The next two lines are the optimal values found through exhaustive search (see text). The last four lines are performance indicators, estimated during the last 100 presentations of each pattern. $\langle P_{\text{learned}} \rangle$ is the mean number of “learned patterns,” that is by convention patterns which elicit at least one postsynaptic spike. The following line is the mean hit rate for those patterns. The subsequent line gives the false alarm rate, but we never observed any here. Finally $P(\text{opt})$ is the proportion of optimal cases.

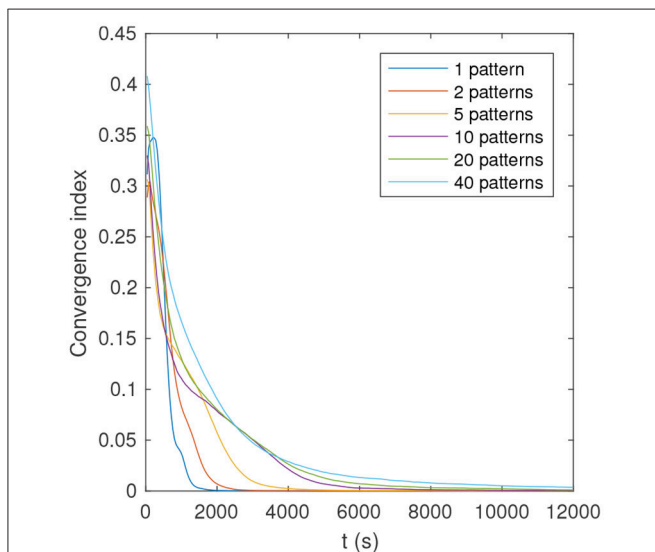


FIGURE 7 | Convergence index as a function of time and number of patterns, for an example of optimal simulation. The convergence index is defined as the mean distance between the full precision weights, and their binary quantization (0 if $w < 0.5$, and 1 otherwise).

optimally selective to several tens of independent patterns. Hence STDP and coincidence detection are compatible with distributed coding.

Yet one issue with having one neuron selective to multiple patterns is stability. If one of the learned pattern does not occur for a long period during which the other patterns occur many times, causing postsynaptic spikes, the unseen pattern will tend to be forgotten. This is not an issue with localist coding: if the learned pattern does not occur, the threshold is hardly ever

reached so the weights are not modified, and the pattern is retained indefinitely, even if STDP is “on” all the time.

Another issue with distributed coding is how the readout is done, that is how the identity of the stimulus can be inferred from multiple neuron responses, given that each response is ambiguous? This is out of the scope of the current paper, but we suspect that STDP could again help. As shown in this study, each neuron equipped with STDP can learn to fire to multiple independent stimuli. Let’s suppose that stimuli are shown one at a time. When stimulus A is shown, all the neurons that learned this stimulus (among others) will fire synchronously. Let us call S this set of neurons. A downstream neuron equipped with STDP could easily become selective to this synchronous volley of spikes from neurons in S (Brette, 2012). With an appropriate threshold, this neuron would fire if and only if all the neurons in S have fired. Does that necessarily mean that A is there? Yes, if the intersection of the sets of stimuli learned by neurons in S only contains A . In the general case, the intersection is likely to be much smaller than the typical sets of stimuli learned by the S neurons, so much of the ambiguity should be resolved.

What could determine the set of patterns to which a neuron responds? Here, we used independent, unrelated, patterns (i.e., with chance-level overlap), and yet several of these patterns could be learned by a single neuron. Of course, patterns with more overlap would be easier to group. So in the presence of multiple postsynaptic neurons, each one would tend to learn a cluster of similar patterns. Another factor is the time at which the patterns are presented: those presented at the same period are more likely to be learned by the same neuron—a neuron which was still unselective at that period. Indeed, neurons equipped with STDP have some sort of critical period, before convergence, during which they can learn new pattern easily. Conversely, after convergence, neurons tend to fire if and only if the patterns they have learned are presented (Figure 6), and thus can hardly learn any new pattern. This is interesting, because patterns presented at the same period are likely to be somewhat related. For example, a neuron could fire to the different people you have met on your first day at work. In the presence of neurogenesis, newborn neurons could handle the learning of other patterns during the subsequent periods of your life. Finally, here we did not use any reward signal. But such a signal, if available, could modulate STDP (leading to some form of supervised learning), and encourage a given neuron to fire to a particular, meaningful, set of patterns (Mozafari et al., 2018a,b), as opposed to a random set like here. For example, a single neuron could learn to fire to any animal, even if different animals cause very different sensory inputs.

Here the STDP rule we used always led to binary weights after learning. That is an afferent could be either selected or discarded. We thus could use our SNR calculations derived with binary weights, and checked that the selected set was optimal given the binary weight constraint. Further calculations in the Appendix (Supplementary Material) suggest that removing such a constraint could lead to a modest increase in SNR , of about 10%. More research is needed to see if a multiplicative STDP rule, which does not converge toward binary weights (van Rossum

et al., 2000; Gütig et al., 2003), could lead to the optimal graded weights.

Our theoretical study suggests, together with others (Gütig and Sompolinsky, 2006; Brette, 2012), that coincidence detection is computationally powerful. In fact, it could be the main function of neurons (Abeles, 1982; König et al., 1996). In line with this proposal, neurons *in vivo* appear to be mainly fluctuation-driven, not mean-driven (Rossant et al., 2011; Brette, 2012, 2015). This is the case in particular in the balanced regime (Brette, 2015), which appears to be the prevalent regime in the brain (Denève and Machens, 2016). Several other points suggest that coincidence detection is the main function of neurons. Firstly, strong feedforward inhibitory circuits throughout the central nervous system often shorten the neurons' effective integration windows (Bruno, 2011). Secondly, the effective integration time constant in dendrites might be one order of magnitude shorter than the soma's one (König et al., 1996). Finally, recent experiments indicate that a neuron's threshold quickly adapts to recent potential values (Platkiewicz and Brette, 2011; Fontaine et al., 2014; Mensi et al., 2016), so that only a sudden potential increase can trigger a postsynaptic spike. This enhances coincidence detection. It remains unclear if other spike time aspects such as ranks (Thorpe and Gautrais, 1998) also matter.

Our results show that lower firing rates lead to better signal-to-noise ratio. It is worth mentioning that mean firing rates are probably largely overestimated in the electrophysiological literature, because extracellular recordings—by far the most popular technique—are totally blind to cells that do not fire at all (Thorpe, 2011). Even a cell that fire only a handful of spikes will be ignored, because spike sorting algorithms need tens of spikes from a given cell before they can create a new cluster corresponding to that cell. Furthermore, experimentalists tend to search for stimuli that elicit strong responses, and, when they can move the electrode(s), tend to look for most responsive cells,

introducing strong selection biases. Mean firing rates, averaged across time and cells, are largely unknown, but they could be smaller than 1 Hz (Shoham et al., 2006). It seems like coding is sparse: neurons only fire when they need to signal an important event, and that every spike matters (Wolfe et al., 2010).

Finally, we see an analogy between our theory, and the one of neural associative memory (NAM), in which an output (data) vector is produced by multiplying an input (address) vector by a weight matrix. Unlike NAM, our framework is dynamic, yet after learning, to a first approximation, our STDP neurons count the number of input spikes arriving through reinforced synapses in a short integration window, and each one outputs a 1 (i.e., a spike) if this count exceeds a threshold, and a 0 otherwise, leading to a binary output vector, much like in a binary NAM. It is thus unsurprising that sparsity is desirable both in our theory, and in NAMs (Palm, 2013).

AUTHOR CONTRIBUTIONS

SK and TM did the analytical derivations. TM performed the numerical simulations. SK and TM analyzed the results. TM wrote the first version of the manuscript. SK reviewed it.

ACKNOWLEDGMENTS

This research received funding from the European Research Council under the European Union's 7th Framework Program (FP/2007-2013) / ERC Grant Agreement n.323711 (M4 project). We thank Milad Mozafari for smart implementation hints, and Jean Pierre Jaffrézou for his excellent copy editing.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fncom.2018.00074/full#supplementary-material>

REFERENCES

- Abeles, M. (1982). Role of the cortical neuron: integrator or coincidence detector? *Isr. J. Med. Sci.* 18, 83–92.
- Abeles, M. (1991). *Corticonics: Neural Circuits of the Cerebral Cortex*. Cambridge: New York, NY: Cambridge University Press.
- Bowers, J. S. (2009). On the biological plausibility of grandmother cells: implications for neural network theories in psychology and neuroscience. *Psychol. Rev.* 116, 220–251. doi: 10.1037/a0014462
- Bowers, J. S. (2017). Parallel distributed processing theory in the age of deep networks. *Trends Cogn. Sci.* 21, 950–961. doi: 10.1016/j.tics.2017.09.013
- Brette, R. (2012). Computing with neural synchrony. *PLOS Comput. Biol.* 8:e1002561. doi: 10.1371/journal.pcbi.1002561
- Brette, R. (2015). Philosophy of the spike: rate-based vs. spike-based theories of the brain. *Front. Syst. Neurosci.* 9:151. doi: 10.3389/fnsys.2015.00151
- Bruno, R. M. (2011). Synchrony in sensation. *Curr. Opin. Neurobiol.* 21, 701–708. doi: 10.1016/j.conb.2011.06.003
- Burkitt, A. N. (2006). A review of the integrate-and-fire neuron model: I. Homogeneous synaptic input. *Biol. Cybern.* 95, 1–19. doi: 10.1007/s00422-006-0068-6
- Denève, S., and Machens, C. K. (2016). Efficient codes and balanced networks. *Nat. Neurosci.* 19, 375–382. doi: 10.1038/nn.4243
- Fontaine, B., Peña, J. L., and Brette, R. (2014). Spike-threshold adaptation predicted by membrane potential dynamics *in vivo*. *PLoS Comput. Biol.* 10:e1003560. doi: 10.1371/journal.pcbi.1003560
- Gerstner, W., and Naud, R. (2009). How good are neuron models? *Science* 326, 379–380. doi: 10.1126/science.1181936
- Gilson, M., Masquelier, T., and Hugues, E. (2011). STDP allows fast rate-modulated coding with Poisson-like spike trains. *PLoS Comput. Biol.* 7:e1002231. doi: 10.1371/journal.pcbi.1002231
- Gütig, R., Aharonov, R., Rotter, S., and Sompolinsky, H. (2003). Learning input correlations through nonlinear temporally asymmetric Hebbian plasticity. *J. Neurosci.* 23, 3697–3714. doi: 10.1523/JNEUROSCI.23-09-03697.2003
- Gütig, R., and Sompolinsky, H. (2006). The tempotron: a neuron that learns spike timing-based decisions. *Nat. Neurosci.* 9, 420–428. doi: 10.1038/nn1643
- Havenith, M. N., Yu, S., Biederlack, J., Chen, N.-H., Singer, W., and Nikolic, D. (2011). Synchrony makes neurons fire in sequence, and stimulus properties determine who is ahead. *J. Neurosci.* 31, 8570–8584. doi: 10.1523/JNEUROSCI.2817-10.2011
- Herrmann, M., Hertz, J. A., and Prügel-Bennett, A. (1995). Analysis of synfire chains. *Netw. Comput. Neural Syst.* 6, 403–414.
- Hines, M. L., Morse, T., Migliore, M., Carnevale, N. T., and Shepherd, G. M. (2004). ModelDB: a database to support computational neuroscience. *J. Comput. Neurosci.* 17, 7–11. doi: 10.1023/B:JCN.0000023869.22017.2e

- Humble, J., Denham, S., and Wennekers, T. (2012). Spatio-temporal pattern recognizers using spiking neurons and spike-timing-dependent plasticity. *Front. Comput. Neurosci.* 6:84. doi: 10.3389/fncom.2012.00084
- Hung, C. P., Kreiman, G., Poggio, T., and DiCarlo, J. J. (2005). Fast readout of object identity from macaque inferior temporal cortex. *Science* 310, 863–866. doi: 10.1126/science.1117593
- Hunzinger, J. F., Chan, V. H., and Froemke, R. C. (2012). Learning complex temporal patterns with resource-dependent spike timing-dependent plasticity. *J. Neurophysiol.* 108, 551–566. doi: 10.1152/jn.01150.2011
- Kasabov, N., Dhoble, K., Nuntalid, N., and Indiveri, G. (2013). Dynamic evolving spiking neural networks for on-line spatio- and spectro-temporal pattern recognition. *Neural Netw.* 41, 188–201. doi: 10.1016/j.neunet.2012.11.014
- Kayser, C., Logothetis, N. K., and Panzeri, S. (2010). Millisecond encoding precision of auditory cortex neurons. *Proc. Natl. Acad. Sci. U.S.A.* 107, 16976–16981. doi: 10.1073/pnas.1012656107
- Kempler, R., Gerstner, W., and van Hemmen, J. L. (1999). Hebbian learning and spiking neurons. *Phys. Rev. E* 59, 4498–4514.
- Kheradpisheh, S. R., Ganjtabesh, M., and Masquelier, T. (2016). Bio-inspired unsupervised learning of visual features leads to robust invariant object recognition. *Neurocomputing* 205, 382–392. doi: 10.1016/j.neucom.2016.04.029
- Kheradpisheh, S. R., Ganjtabesh, M., Thorpe, S. J., and Masquelier, T. (2018). STDP-based spiking deep convolutional neural networks for object recognition. *Neural Netw.* 99, 56–67. doi: 10.1016/j.neunet.2017.12.005
- Klampfl, S., and Maass, W. (2013). Emergence of dynamic memory traces in cortical microcircuit models through STDP. *J. Neurosci.* 33, 11515–11529. doi: 10.1523/JNEUROSCI.5044-12.2013
- Kobayashi, R., Tsubo, Y., and Shinomoto, S. (2009). Made-to-order spiking neuron model equipped with a multi-timescale adaptive threshold. *Front. Comput. Neurosci.* 3:9. doi: 10.3389/fncom.2009.10.009.2009
- König, P., Engel, A. K., and Singer, W. (1996). Integrator or coincidence detector? The role of the cortical neuron revisited. *Trends Neurosci.* 19, 130–137.
- Krunglevicius, D. (2015). Competitive STDP learning of overlapping spatial patterns. *Neural Comput.* 27, 1673–1685. doi: 10.1162/NECO_a_00753
- LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature* 521, 436–444. doi: 10.1038/nature14539
- Masquelier, T. (2013). Neural variability, or lack thereof. *Front. Comput. Neurosci.* 7, 1–7. doi: 10.3389/fncom.2013.00007
- Masquelier, T. (2017). STDP allows close-to-optimal spatiotemporal spike pattern detection by single coincidence detector neurons. *Neuroscience* doi: 10.1016/j.neuroscience.2017.06.032. [Epub ahead of print].
- Masquelier, T., Guyonneau, R., and Thorpe, S. J. (2008). Spike timing dependent plasticity finds the start of repeating patterns in continuous spike trains. *PLoS ONE* 3:e1377. doi: 10.1371/journal.pone.0001377
- Masquelier, T., Guyonneau, R., and Thorpe, S. J. (2009). Competitive STDP-based spike pattern learning. *Neural Comput.* 21, 1259–1276. doi: 10.1162/neco.2008.06-08-804
- Masquelier, T., and Thorpe, S. J. (2007). Unsupervised learning of visual features through spike timing dependent plasticity. *PLoS Comput. Biol.* 3:e31. doi: 10.1371/journal.pcbi.0030031
- Mensi, S., Hagens, O., Gerstner, W., and Pozzorini, C. (2016). Enhanced sensitivity to rapid input fluctuations by nonlinear threshold dynamics in neocortical pyramidal neurons. *PLOS Comput. Biol.* 12:e1004761. doi: 10.1371/journal.pcbi.1004761
- Mozafari, M., Ganjtabesh, M., Nowzari-Dalini, A., Thorpe, S. J., and Masquelier, T. (2018a). Combining STDP and reward-modulated STDP in deep convolutional spiking neural networks for digit recognition. *arXiv[preprint]* arXiv:1804.00227.
- Mozafari, M., Kheradpisheh, S., Masquelier, T., Nowzari-Dalini, A., and Ganjtabesh, M. (2018b). First-spike-based visual categorization using reward-modulated STDP. *IEEE Trans. Neural Netw. Learn. Syst.* doi: 10.1109/TNNLS.2018.2826721. [Epub ahead of print].
- Nessler, B., Pfeiffer, M., Buesing, L., and Maass, W. (2013). Bayesian computation emerges in generic cortical microcircuits through spike-timing-dependent plasticity. *PLoS Comput. Biol.* 9:e1003037. doi: 10.1371/journal.pcbi.1003037
- Nguyen, A., Dosovitskiy, A., Yosinski, J., Brox, T., and Clune, J. (2016). “Synthesizing the preferred inputs for neurons in neural networks via deep generator networks,” in *Advances in Neural Information Processing Systems*.
- Olah, C., Mordvintsev, A., and Schubert, L. (2017). Feature visualization. *Distill*. doi: 10.23915/distill.00007
- O’Reilly, R. C. (1998). Six principles for biologically based computational models of cortical cognition. *Trends Cogn. Sci.* 2, 455–462.
- Palm, G. (2013). Neural associative memories and sparse coding. *Neural Netw.* 37, 165–171. doi: 10.1016/j.neunet.2012.08.013
- Panzeri, S., and Diamond, M. E. (2010). Information carried by population spike times in the whisker sensory cortex can be decoded without knowledge of stimulus time. *Front. Synapt. Neurosci.* 2:17. doi: 10.3389/fnsyn.2010.00017
- Platkiewicz, J., and Brette, R. (2011). Impact of fast sodium channel inactivation on spike threshold dynamics and synaptic integration. *PLoS Comput. Biol.* 7:e1001129. doi: 10.1371/journal.pcbi.1001129
- Quian Quiroga, R., and Kreiman, G. (2010). Measuring sparseness in the brain: comment on bowers (2009). *Psychol. Rev.* 117, 291–297. doi: 10.1037/a0016917
- Quiroga, R. Q., Kreiman, G., Koch, C., and Fried, I. (2008). Sparse but not “Grandmother-cell” coding in the medial temporal lobe. *Trends Cogn. Sci.* 12, 87–91. doi: 10.1016/j.tics.2007.12.003
- Rolls, E. T., Treves, A., and Tovee, M. J. (1997). The representational capacity of the distributed encoding of information provided by populations of neurons in primate temporal visual cortex. *Exp. Brain Res.* 114, 149–162.
- Rossant, C., Leijon, S., Magnusson, A. K., and Brette, R. (2011). Sensitivity of noisy neurons to coincident inputs. *J. Neurosci.* 31, 17193–17206. doi: 10.1523/JNEUROSCI.2482-11.2011
- Roy, A. (2017). The theory of localist representation and of a purely abstract cognitive system: the evidence from cortical columns, category cells, and multisensory neurons. *Front. Psychol.* 8:186. doi: 10.3389/fpsyg.2017.00186
- Shoham, S., O’Connor, D. H., and Segev, R. (2006). How silent is the brain: is there a “dark matter” problem in neuroscience? *J. Compar. Physiol.* 192, 777–784. doi: 10.1007/s00359-006-0117-6
- Song, S., Miller, K. D., and Abbott, L. F. (2000). Competitive hebbian learning through spike-timing-dependent synaptic plasticity. *Nat. Neurosci.* 3, 919–926. doi: 10.1038/78829
- Sun, H., Sourina, O., and Huang, G.-B. (2016). Learning polychronous neuronal groups using joint weight-delay spike-timing-dependent plasticity. *Neural Comput.* 28, 2181–2212. doi: 10.1162/NECO_a_00879
- Thorpe, S. (1989). Local vs. distributed coding. *Intellectica* 8, 3–40.
- Thorpe, S. J. (2009). Single units and sensation: still just as relevant today. *Perception* 38, 802–803; discussion: 804–807.
- Thorpe, S. J. (2011). “Grandmother cells and distributed representations,” in *Visual Population Codes-Toward a Common Multivariate Framework for Cell Recording and Functional Imaging*, eds N. Kriegeskorte, G. Kreiman (Cambridge, MA: MIT Press), 23–51.
- Thorpe, S. J., and Gautrais, J. (1998). “Rank order coding,” in *Computational Neuroscience: Trends in Research*, ed J. M. Bower (New York, NY: Plenum Press), 113–118.
- van Rossum, M. C., Bi, G. Q., and Turrigiano, G. G. (2000). Stable Hebbian learning from spike timing-dependent plasticity. *J. Neurosci.* 20, 8812–8821. doi: 10.1523/JNEUROSCI.20-23-08812.2000
- Wolfe, J., Houweling, A. R., and Brecht, M. (2010). Sparse and powerful cortical spikes. *Curr. Opin. Neurobiol.* 20, 306–312. doi: 10.1016/j.conb.2010.03.006
- Yger, P., Stimberg, M., and Brette, R. (2015). Fast learning with weak synaptic plasticity. *J. Neurosci.* 35, 13351–13362. doi: 10.1523/JNEUROSCI.0607-15.2015
- Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., and Torralba, A. (2015). Object detectors emerge in deep scene CNNs. *arXiv[preprint]* arXiv:1412.6856.

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2018 Masquelier and Kheradpisheh. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.