Please cite this article in press as: Masquelier T. STDP allows close-to-optimal spatiotemporal spike pattern detection by single-coincidence detector neurons. Neuroscience (2017), http://dx.doi.org/10.1016/j.neuroscience.2017.06.032

Neuroscience xxx (2017) xxx–xxx

STDP ALLOWS CLOSE-TO-OPTIMAL SPATIOTEMPORAL SPIKE PATTERN DETECTION BY SINGLE-COINCIDENCE DETECTOR NEURONS

TIMOTHÉE MASQUELIER
CERCO UMR5549 CNRS – Université Toulouse 3, France

Abstract—Repeating spatiotemporal spike patterns exist and carry information. How this information is extracted by downstream neurons is unclear. Here we theoretically investigate to what extent a single cell could detect a given spike pattern and what are the optimal parameters to do so, in particular the membrane time constant τ. Using a leaky integrate-and-fire (LIF) neuron with homogeneous Poisson's input, we computed this optimum analytically. We found that a relatively small τ (at most a few tens of ms) is usually optimal, even when the pattern is much longer. This is somewhat counter-intuitive as the resulting detector ignores most of the pattern, due to its fast memory decay. Next, we wondered if spike-timing-dependent plasticity (STDP) could enable a neuron to reach the theoretical optimum. We simulated a LIF equipped with additive STDP, and repeatedly exposed it to a given input spike pattern. As in previous studies, the LIF progressively became selective to the repeating pattern with no supervision, even when the pattern was embedded in Poisson's activity. Here we show that, using certain STDP parameters, the resulting pattern detector is optimal. These mechanisms may explain how humans learn repeating sensory sequences. Long sequences could be recognized thanks to coincidence detectors working at a much shorter timescale. This is consistent with the fact that recognition is still possible if a sound sequence is compressed, played backward, or scrambled using 10-ms bins. Coincidence detection is a simple yet powerful mechanism, which could be the main function of neurons in the brain.

This article is part of a Special Issue entitled: Sequence Processing

Key words: spike-timing-dependent plasticity (STDP), leaky integrate-and-fire neuron, coincidence detection, multi-neuron spike sequence, spatiotemporal spike pattern, unsupervised learning.

INTRODUCTION

Electrophysiologists report the existence of repeating spike sequence involving multiple cells, also called “spatiotemporal spike patterns”, with precision in the millisecond range, both *in vitro* and *in vivo*, lasting from a few tens of ms to several seconds (Tiesinga et al., 2008). In sensory systems, different stimuli evoke different spike patterns (also called “packets”) (Luczak et al., 2015). In other words, the spike patterns contain information about the stimulus. How this information is extracted by downstream neurons is unclear. Can it be done by neurons only one synapse away from the recorded neurons? Or are multiple integration steps needed? Can it be done by simple coincidence detector neurons, or should other temporal features, such as spike ranks, be taken into account? Here we wondered how far we can go with the simplest scenario: the readout is done by simple coincidence detector neurons only one synapse away from the neurons involved in the repeating pattern. We demonstrate that this approach can lead to very robust pattern detectors, provided that the membrane time constants are relatively short, possibly much shorter than the pattern duration.

In addition, it is known that mere repeated exposure to meaningless sensory sequences facilitates their recognition afterward, in the visual (Gold et al., 2014) and auditory modalities (Agus et al., 2010; Andrillon et al., 2015; Viswanathan et al., 2016) (see also contributions in this special issue), even when the subjects were unaware of these repetitions. Thus, an unsupervised learning mechanism must be at work. It could be the so called spike-timing-dependent plasticity (STDP). Indeed, some theoretical studies by us and others have shown that neurons equipped with STDP can become selective to arbitrary repeating spike patterns, even without supervision (Masquelier et al., 2008, 2009; Gilson et al., 2011; Humble et al., 2012; Hunzinger et al., 2012; Klampfl and Maass, 2013; Kasabov et al., 2013; Nessler et al., 2013; Krunglevicius, 2015; Yger et al., 2015; Sun et al., 2016). Using numerical simulations, we show here that the resulting detectors can be close to the theoretical optimum.

FORMAL DESCRIPTION OF THE PROBLEM

We assess the problem of detecting a spatiotemporal spike pattern with a single-LIF neuron. Intuitively, one should connect the LIF to the neurons that are particularly active during the pattern, or during a subsection of it. That way, the LIF will tend to be more activated by the pattern than by some other input. More formally, we note L the pattern duration, N the number of neurons it involves. We call Strategy $\#n$ the strategy

---

Abbreviations: LTD, Long Term Depression; LIF, leaky integrate-and-fire; SNR, signal-to-noise ratio; STDP, spike-timing-dependent plasticity.

E-mail address: timothee.masquelier@cnrs.fr

http://dx.doi.org/10.1016/j.neuroscience.2017.06.032
0306-4522/

which consists in connecting the LIF to the M neurons that emit at least n spike(s) during a certain time window $\Delta t \leqslant L$ of the pattern. Strategy #1 is illustrated on Fig. 1.

We hypothesize that all afferent neurons fire according to an homogeneous Poisson process with rate f, both inside and outside the pattern. That is the pattern corresponds to one realization of the Poisson process, which can be repeated (this is sometimes referred to a "frozen noise"). To model jitter, at each repetition a random time lag 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, see next section).

We also assume that synapses are instantaneous (i.e. excitatory postsynaptic currents are made of Diracs), which facilitate the analytic calculations.

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

$$SNR = \frac{V_{\max} - \overline{V}_{\text{noise}}}{\sigma_{\text{noise}}}, \tag{1}$$

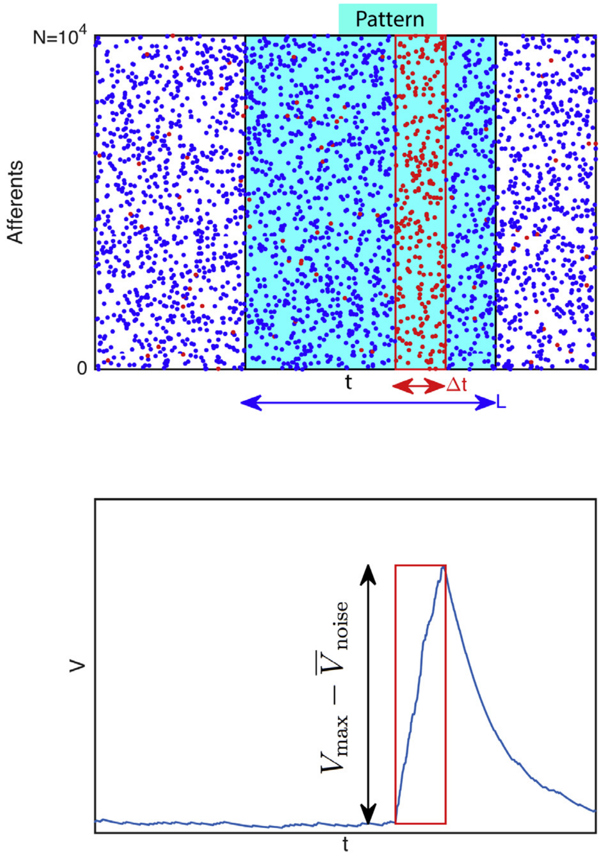

Fig. 1. Detecting a spike pattern with a LIF neuron. (Top) Raster plot of $N = 10^4$ neurons firing according to an homogeneous Poisson process. A pattern of duration L can be repeated (frozen noise). Here we illustrated Strategy #1, which consists in connecting the LIF to all neurons that fire at least once during a certain time window of the pattern, with duration $\Delta t \leqslant L$. These neurons emit red spikes. Of course they also fire outside of the Δt window. (Bottom) Typically the LIF's potential will be particularly high when integrating the spikes of the Δt window, much higher than with random Poisson's inputs, and we want to optimize this difference, or more precisely the signal-to-noise ratio (SNR, see text).

where $V_{\max}$ is the maximal potential reached during the pattern presentation, $\overline{V}_{\text{noise}}$ is the mean value for the potential with Poisson's input (noise period), and σ_{noise} its standard deviation (see Fig. 1).

A THEORETICAL OPTIMUM

Deriving the SNR analytically

We now want to calculate the SNR analytically. In this section, we assume unitary synaptic weights. Since the LIF has instantaneous synapses, and the input spikes are generated with a Poisson process, we have $\overline{V}_{\text{noise}} = \tau f M$ and $\sigma_{\text{noise}} = \sqrt{\tau f M / 2}$, where τ is the membrane's time constant (Burkitt, 2006). We assume that $\tau f M \gg 1$ (large number of synaptic inputs), so that the distribution of V is approximately Gaussian (Burkitt, 2006). Otherwise it would be positively skewed, thus a high SNR as defined by Eq. (1) would not guarantee a low false alarm rate.

The number of selected afferents M depends on the strategy n. The probability that an afferent fires k times in the Δt window is given by the Poisson probability mass function: $\mathrm{P}(k\text{spikes}) = \frac{\lambda^k e^{-\lambda}}{k!}$, with $\lambda = f\Delta t$. The probability that an afferent fires at least n times is thus $1 - e^{-\lambda}\sum_{k=0}^{n-1}\frac{\lambda^k}{k!}$, and finally, on average:

$$M = N\left(1 - e^{-\lambda}\sum_{k=0}^{n-1}\frac{\lambda^k}{k!}\right). \tag{2}$$

We now need to estimate $V_{\max}$. Intuitively, during the Δt window, the effective input spike rate, which we call r, is typically higher than fM, because we deliberately chose the most active afferents. For example, using Strategy #1 with $\Delta t = 10$ ms ensures that this rate is at least 100 Hz per afferent, even if f is only a few Hz. More formally, Strategy #n discards the afferents that emit fewer than n spikes. This means on average the number of discarded spikes is $Ne^{-\lambda}\sum_{k=0}^{n-1}\frac{k\lambda^k}{k!} = Ne^{-\lambda}\sum_{k=1}^{n-1}\frac{\lambda^k}{(k-1)!} = Ne^{-\lambda}\lambda\sum_{k=1}^{n-1}\frac{\lambda^{k-1}}{(k-1)!} = Ne^{-\lambda}\lambda\sum_{k=0}^{n-2}\frac{\lambda^k}{k!}$. Thus on average:

$$r = N/\Delta t\left(\lambda - e^{-\lambda}\lambda\sum_{k=0}^{n-2}\frac{\lambda^k}{k!}\right) = Nf\left(1 - e^{-\lambda}\sum_{k=0}^{n-2}\frac{\lambda^k}{k!}\right). \tag{3}$$

We note $\overline{V}^{\infty} = \tau r$ the mean potential of the steady regime that would be reached if Δt was infinite. We now want to compute the transient response. The LIF with instantaneous synapses and unitary synaptic weights obeys the following differential equation:

$$\tau\frac{\mathrm{d}V}{\mathrm{d}t} = -V + \tau\sum_i \delta(t - t_i), \tag{4}$$

where t_i are the presynaptic spike times. We first make the approximation of continuity, and replace the sum of Diracs by an equivalent firing rate $R(t)$:

$$\tau\frac{\mathrm{d}V}{\mathrm{d}t} = -V + \tau R(t). \tag{5}$$

$R(t)$ should be computed on a time bin which is much smaller than τ, but yet contains many spikes, to avoid discretization effects. In other words, this approximation of continuity is only valid for a large number of spikes in

Please cite this article in press as: Masquelier T. STDP allows close-to-optimal spatiotemporal spike pattern detection by single-coincidence detector neurons. Neuroscience (2017), http://dx.doi.org/10.1016/j.neuroscience.2017.06.032