

A critical survey of STDP in Spiking Neural Networks for Pattern Recognition

Alex Vigneron

Univ. Lille, CNRS, Centrale Lille, UMR 9189 – CRISTAL
F-59000 Lille, France
a.vigneron@protonmail.com

Jean Martinet

Université Côte d’Azur, CNRS, I3S
France
jean.martinet@univ-cotedazur.fr

Abstract—The bio-inspired concept of Spike-Timing-Dependent Plasticity (STDP) derived from neurobiology is increasingly used in Spiking Neural Networks (SNNs) nowadays. Mostly found in unsupervised learning, though recent work has shown its usefulness in supervised or reinforced paradigms too, STDP is a key element to understanding SNN architectures’ learning process. This review introduces a categorisation of its several variants and discusses their specificities and applications, from a pattern recognition perspective. It gathers a variety of definitions used in machine learning for pattern recognition. It provides relevant information for research communities of various backgrounds looking for an overview of this field.

Index Terms—Spiking Neural Networks, Machine Learning, Artificial Neural Networks, Pattern Recognition, Unsupervised Learning, STDP, Bio-inspiration.

INTRODUCTION

Machine learning and pattern recognition domains are vastly dominated by deep learning. In less than a decade, deep artificial neural networks (based on formal neurons) have successfully pulled state-of-the-art performances of machine learning tasks to new levels, on a wide range of challenging benchmarks. The availability of both tremendous amounts of annotated data and huge computational resources have enabled remarkable progress. However, this success comes with substantial costs in both human intervention for data labelling and energy for training, despite most recent advances in parallel digital architectures. Regarding data, for instance, one of the top-ranking systems in LFW, the major face recognition challenge, is *IFLYTEK-CV* (currently ranked third¹), was trained using a dataset of 3.8M face images of 85K individuals. This amount, which has now become standard, is far beyond what was usual just a decade ago. Because such methods heavily rely on stochastic gradient descent and back-propagation, they require tremendous computational power. For instance, ResNet [11] has been trained for 3 weeks on a 8-GPU server, which is equivalent to a power consumption of about 1 GWh. More generally, worldwide data centres currently require a power of about 1 PW, equivalent to 4% of GHG emissions, which exceeds those of air transportation.

This work has been partly funded by IRCICA (Univ. Lille, CNRS, USR 3380 IRCICA, F-59000 Lille, France) under the Bioinspired Project. This work was also supported by the French government through the Program “Investissements d’avenir” (I-ULNE SITE / ANR-16-IDEX-0004 ULNE) managed by the National Research Agency.

¹See vis-www.cs.umass.edu/lfw/results.html. Last accessed May 2020.

Forecasts plan that this figure will double every 4 years. Hence, a paradigm change in machine learning and pattern recognition is needed in order to face the ever-growing demand.

Spiking Neural Networks display promising characteristics for this paradigm change [23], [25], [30], [34], such as unsupervised training with STDP rules, which reduces the need for large annotated datasets. SNNs show higher efficiency than classical neural networks, from both computation [18] and energy [8] points of view. First, regarding computation, with temporal coding (see Section I-C), the core information in SNN models lies in the very timing of binary spikes, and do not require to manipulate large matrices of floating-point numbers. Second, the model is intrinsically sparse since units only fire when needed. Therefore, just a few (thousands of) spikes are needed at inference, as opposed to a total network activation with classical neurons. Finally, the STDP rule for a given synapse only involves the local spike timings of the pre- and the post- synaptic neurons. This locality of the STDP rule makes it hardware-friendly, contrary to stochastic gradient descent, which requires a global loss differentiation. Hence, SNNs are highly energy efficient when implemented on neuromorphic hardware. Ultra-low-power neuromorphic hardware implementing SNNs can be built with CMOS technology, and typically uses below-threshold voltage, enabling to reduce energy dissipation by several orders of magnitude, compared to standard digital architectures, even when using special-purpose accelerator hardware such as Tensor Processing Units.

Historically, STDP started with Hebb’s remark that “*when an axon of cell A is near enough to excite a cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A’s efficiency, as one of the cells firing B, is increased*” [12]. Hebb consistently insists on the causality between afferent \mathcal{N}_{aff} firing efferent \mathcal{N}_{eff} stating that “*one cell repeatedly assists in firing another*”. Therefore appeared as an attractive option since it followed neurobiological processes governing the brain. Furthermore, the brain being a power-efficient computing unit, biomimeticism implied increased power efficiency which was the second motivation behind such enthusiasm [15].

STDP learning rule, implementing Hebbian learning, adjusts synaptic strength through correlation detection between pre- and post- synaptic firing times therefore affecting performances [26]. Its mechanism is inspired from the organic

hierarchical ventral path of the visual cortex where neurons communicate only through discrete spikes and not continuous information [15]. Despite literature converging towards the idea that most of the learning is performed in unsupervised ways, some feedback and reinforcement connections and techniques are used in the brain, and it seems that purely unsupervised approach in the biological brain only serves as a basic tool of image recognition before the more fine-grained cogs come into action [15], [20]. The supposed secondary role of feedback in image recognition is grounded in the biological findings that the primate’s visual cortex extracts basic information from an image in about 100 ms [15], thus supporting the theory that at least part of the process is feasible in a pure feed-forward way.

Biological STDP has been shown to produce both Long-Term Potentiation (LTP) in synapses when the pre-synaptic neuron had fired shortly before the post-synaptic one and Long-Term Depression (LTD) when the pre-synaptic neuron had fired shortly after the post-synaptic one. Synapses which \mathcal{N}_{affs} had not fired either before or after the \mathcal{N}_{effs} are left untouched. The modification of neural potential can occur following several variants of STDP. Synaptic potential can be tuned either through additive STDP where any correlated inputs will be added to the potential of the post-synaptic neuron, no matter how distant from firing time, and the multiplicative STDP rule where a coefficient based on time-remoteness from the signal is taken into account. A brief overview of the several variants of STDP used in SNN is presented in I together with their current examples of use.

I. SPIKING NEURAL NETWORKS

We consider both the afferent neuron (abbreviated *aff*, also termed *pre-synaptic*) and the efferent neuron (abbreviated *eff*, also termed *post-synaptic*) on each side of a given synapse. We define $\delta = t_{eff} - t_{aff}$ as the time coefficient linked to LTP if $\delta \geq 0$, LTD otherwise, where t_{eff} (resp. t_{aff}) is the spike timing of the efferent (resp. afferent) neuron.

A. Neurons

Despite the wide range of existing artificial neurons, only a few popular models seem to draw most of the attention, namely IF and LIF neurons, because of their simplified equations and relative biological inspiration. In the specific case of the LIF \mathcal{N} , the role of the τ -leak parameter is to replicate biological models’ preparation for the next spike, however an alternative biologically plausible solution would be the systematic reset of membrane potential to a resting level between spikes [19]. Changing the neuron’s model does not necessarily imply wide variations in performance, hence inferring that computation is robust to increasing complexity of computational unit [32], though this phenomenon can be more or less common depending on STDP rule. Dropout techniques used in training consist in ignoring several neurons at random. These improve the global involvement rate of neurons while concurrently decreasing the networks’ complexity [20].

B. Synapses and synaptic weights ω

The neurobiological synapse \mathcal{S} -synaptic cleft in the case of chemical synapses- is a region between \mathcal{N}_{aff} and \mathcal{N}_{eff} which acts as a communication interface between both. In biology, whether they interact directly through electrical action potential or through neurotransmitters modifying cytoplasmic content which in turn cause action potential, those communication paradigms are predominantly local and quick (tens of milliseconds). According to Hebb’s theory, synapses which often interact develop some “*growth process or metabolic change*” [12] therefore increasing the likelihood of future collaborations between \mathcal{N}_{aff} and \mathcal{N}_{eff} . In computer science, this scheme has been modelled by an \mathcal{S} object which main attribute is its synaptic weight ω , corresponding to Hebb’s *growth*. Minimal efficiency corresponds to a disconnected synapse where firing of an \mathcal{N}_{aff} does not increase the \mathcal{N}_{eff} ’s membrane potential. [32]. Initialisation of ω is often set according to a random [19] normal distribution with $\omega \rightsquigarrow \mathcal{N}(\mu, \sigma)$ where μ should not be too small in order to avoid *dead* neurons (i.e. never firing since their thresholds would never be reached) nor σ too large since it would increase dependency on initial random values with some synapses overpowering others in terms of contribution [15]. In the case of R-STDP, small initial μ values result in hard to train neurons while large σ values result in increased impact of the initial random distribution, henceforth the choice of high μ and low σ is considered optimal [20]. These synaptic weights vary during the learning process and should stabilise upon reaching an equilibrium point [32] yet usually need to be paired with a decaying term to avoid unbounded results [19], except in the case of P-STDP [32]. Their stability greatly depends on the type of STDP used. In computer vision applications, monitoring the dynamic of those weights can allow retracing of implicit representation of models [21] making it possible to reconstruct archetypal images from learned features [9]. This and the observation that these synaptic weights’ distributions are often bi-modal after training with weights being either close to 0 or 1 [9] led to research trying to modify the learning rule in order to allow weights to adopt a more stable unimodal distribution and represent more fine-tuned features [22]. Synaptic weights’ dynamics are determined by the learning rate hyper-parameter α . As $\alpha \rightarrow +\infty$, learning memory decreases making the network *forget* previous images faster; on the other hand, as $\alpha \rightarrow -\infty$, the learning process slows down. Also, the ratio between α^+ and α^- must take into account that $\mathcal{P}(\text{synapse undergoing LTD}) > \mathcal{P}(\text{synapse undergoing LTP})$ -particularly at the beginning- and should therefore remain positive with $\alpha^+ > \alpha^-$ yet not overwhelmingly so [15]. It has been shown that it is possible to generate networks tolerant to synaptic variability which are robust against learning rate and weight initialisation values, therefore avoiding fine hyper-parameter tuning [1].

C. Information coding

This mechanism is inspired by the organic hierarchical ventral path of the visual cortex where neurons communicate

STDP variants and their specificities			
STDP type	Notable features	Limitations	Application examples
C_a -STDP	simple implementation robust to noise tolerant to synaptic variability	unstable skewed distribution of ω limited resistance to jitter poor differentiation between highly similar objects scarce support of intra-class variance	feature extraction for classification (Masquelier, 2007) [19]
C_m -STDP	semi-stable based on δ and ω more adaptative than C_a -STDP	skewed distribution of ω poor differentiation between highly similar objects	colour image recognition (Falez, 2019) [9] : 48.27% (CIFAR-10) — 25.20% (CIFAR-100) — 49.20% (STL-10) large-scale image recognition (Kheradpisheh, 2018) [15] : 99.1% (Caltech) — 98.4% (MNIST) — 82.2% (ETH-80) trajectory detection, AER (Bichler, 2012) [1] : 98% (trajectory) — 95% (counting) grey-scale image reco. (Nessler, 2009) [21]
M-STDP	high LTP correlation supports feed- forward and backward	limited biological plausibility	reproduction of AE learning (Burbank, 2015) [5]
P-STDP	simple implementation (based on δ) robust	unstable skewed distribution of ω	grey-scale image classification (Tavanaei, 2016) [32] : 99% (Caltech-face) — 97.5%(Caltech-motorbike)
R-STDP	high LTP correlation focus on discriminative features	unstable skewed distribution of ω	grey-scale image classification (Mozafari, 2018) [20] : 98.9%(Caltech) — 89.5%(ETH-80) — 88.4%(NORB)
Rev-STDP	stable top-down ω in the case of depression-biased learning possibility to combine with bottom-up variants	high focus on correlation	numerical simulations (Burbank, 2012) [6]
S-STDP	normal ω distribution without necessity of explicit bounds	more complex than C_m -STDP from which it evolved	optical flow, AER data (Paredes-Vallès, 2019) [22]
T-STDP	stable ω supports dense/overlapping time-windows	complex time-window definition	differentiating mutually inclusive spatial patterns (Kruglevicius, 2016) [17]

TABLE I
A COMPARATIVE OVERVIEW OF STDP LEARNING RULES

This table summarises all variations of STDP discussed in this paper. Both C-STDPs are the most basic rules that apply LTP/LTD according to the presence/absence of neural connection. M-STDP’s particularity lies in its correlation time window being centered on the efferent neuron. For P-STDP, variations in ω are exponentially related to current ω , aiming for no change at convergence. R-STDP in turn relies on reinforcement learning approach, regulating synaptic behaviour with reward signals. S-STDP introduces a parameter to monitor excitability. Rev-STDP simulates biological top-down communication by reversing the roles of efferent and afferent neurons. Finally, T-STDP relies on three neurons instead of two.

only through discrete spikes without continuous information [15]. Biological spikes are brief (1 or 2 milliseconds) discrete events, the way in which they encode information in living organisms varies with the type of stimulus and neurons involved. Both frequency and temporal coding rely on spikes through time to convey information, yet their approaches diverge on *what* carries information, whether spike frequency or timing or both. Since correlation and co-variation can not be directly assimilated to causality [4] the question has not yet been tipped in favour of neither one nor the other. Amplitude and duration of spikes don’t vary much, which implies that the semantic content must lie in timing [15]. Neurobiological approaches suggest rate coding alone fails to account for the speed of data transfer, hence pointing towards a, at least partially, temporal coding [17]. Another argument in disfavour of frequency coding is its energy cost due to the high number of spikes needed to encode information [15]. Both evoked reasons tip the scales in temporal coding’s favour which quickens the process with more strongly activated neurons firing earlier [19].

D. Lateral inhibition

The use of STDP in WTA setting can be interpreted as an implementation of Expectation-Maximisation for retracing the causes of neuronal inputs [21]. This competition mechanism consists in having the target \mathcal{N}_{aff} send a signal to other neurons from the same layer in order to prevent them from firing [9]. It was observed as part of cortical architecture and implemented in ASNN to enhance network performance and encourage selectivity. Indeed, the main goal of lateral inhibition is to prevent too many efferent neurons to receive too many spikes from an afferent particular region, thus ensuring that different neurons code different inputs [24]. In the brain, lateral inhibition leads to WTA situations, the most commonly used being k-WTA rules, yet other forms exist such as inter-group WTA competition instead of intra-group, leading to a cohort of winner-neurons rather than a single one. This latter approach falling outside the scope of this extended abstract, we refer to Xie et al.’s [36] for a detailed analysis of the dynamics of such system. We shall here focus on staple neuron to neuron competition, leading to k-WTA situations (where often $k = 1$). In this paradigm, only \mathcal{N} neurons are allowed to spike per either image or frame when

subdivided into local rule [9]. Despite WTA's core purpose being enhanced selectivity, it also increases feature sparsity [9] and maintains synaptic weights within bounds [19]. In terms of configuration, two aspects are at stake, namely choosing the stage up to which lateral inhibition should be maintained and determining its range of action. For a detailed discussion on the latter, see Rolls and Milward's work [24]. When it comes to the former, it has been observed that once the neurons are specialised, lateral inhibition should be discarded since it could prevent target neurons to fire at the right stimuli because of temporally overlapping features [1]. Moreover, as neurons become more selective, lateral inhibition loses its importance and becomes redundant once learning is stabilised as long as STDP is deactivated too [1]. Lateral inhibition's range of action should be correlated to the afferent-efferent wiring scheme, henceforth a local lateral inhibition rule with fully connected layers would make little sense and allow multiple \mathcal{N}_{eff} to spike with the same pattern. Experiments showed that the best ratio for lateral inhibition radius within a layer was approximately half of the breadth of the excitatory connections [24]. Authors [9] suggest that the classical WTA inhibition rule might be detrimental to the network since despite producing sparsity it does not prevent neurons from adapting to similar and therefore redundant features. In some cases though, WTA can cause several output neurons to over-specialise [21]. Some implementations use lateral inhibition proportionally to the distance from the spike with a linear decrease [1].

E. Network architectures

The most common structure is a feed-forward one with varying degrees of connectivity between layers, the simplest one being full connection. Bio-inspired SNNs are for now mostly shallow networks since the transmission of synaptic weights is still unstable. Networks usually consist of at least three layers among which only the middle one is made of SNNs. The initial pre-processing layer transforms pixel values into spike-trains while the final supervised layer performs classification with one to several pooling and convolutional layers in between. Though network architectures impact learning performances of networks, Sboev et al. [26] demonstrated that the type of STDP used had a major impact on them too.

II. STDP LEARNING RULES

STDP is often considered as the implementation of Hebb's synaptic plasticity theory [22] boiling down to synapses being either reinforced or depressed according to δ [15] and can be modelled as a function of δ [17], combining simple features into complex patterns based on statistics [19]. In other words, STDP can be seen as a Hebbian rule in the temporal domain [19] acting as a coincidence and correlation detector [1], [20]. It mainly consists in augmenting synaptic weights of afferent neurons involved in efferent neuron's firing [19], though some exceptions such as M-STDP are found. The variety of STDP rules developed in computer science should come as no surprise since in biological settings, several STDP rules exist depending on the type of synapse (excitatory,

inhibitory) they interact with [17], [28]. For instance, the inversion of LTP and LTD on synapses from the same neuron can occur in biological settings depending on the distance separating them from the soma [17]. The final goal of STDP being to develop enough selectivity so as to minimise false positives and maximise invariance in order to minimise missed images [19]. Moreover, the process should be fast to preserve neurobiological plausibility since processing of visual signal in the brain takes at most 100ms [19].

A. Classical-STDP (C-STDP)

Both additive C_a -STDP and multiplicative C_m -STDP fall within the scope of what we label as C-STDP. On the one hand, additive STDP relies exclusively on δ and is therefore unstable, requiring extra constraints in order to bind synaptic weights which will ultimately result in a skewed distribution with virtually all weights gathered around superior and inferior bounds [22]. For further simplification, some instances of C-STDP only use δ and not its value [1]. On the other hand, multiplicative STDP is slightly more complex since it relies on δ with an added notion of proportionality in terms of weights [22]. In this case also termed *soft bounds* the larger the synaptic weight, the more LTD increases in relation to LTP [28]. Yet despite being more adaptive this rule also results in a distribution of weights skewed towards the bounds. In some variants [1], all the synapses which had not participated directly in the post-synaptic potentiation were depressed while those which had participated were potentiated. Applying LTD to synapses which had not been activated at all (neither before nor after post-synaptic neuron) contrasts with biological STDP rules where only activated yet wrongly-timed synapses receive a depression.

B. Mirror-STDP (M-STDP)

M-STDP's main postulate is that the time window to be analysed for correlation between afferent and efferent spikes is to be centred on the efferent spike, opening the possibility to correlating $\delta < 0$ to LTP instead of LTD. This relies on a simplification of complex unknown mechanisms that systematically cause \mathcal{N}_{eff} and \mathcal{N}_{aff} to fire *together*, its biological plausibility is unsatisfactory. The main advantage of this implementation being that it brings together feed-forward and feed-backward paradigms [5]. This STDP neglects the causality underlined by Hebb [12] with "*takes part in firing*" to instead account for correlated mechanisms with no causality. For a detailed review on biological mechanisms of STDP variations according to the nature of synapses - excitatory/inhibitory- involved and why it is unlikely that $\delta < 0$ could induce LTP in biological setting for an excitatory synapse, see Caporale et al. [7].

C. Probabilistic-STDP (P-STDP)

P-STDP was introduced by Tavanaei et al. [32] in 2016, inspired by Masquelier and Thorpe [19]. With this rule, all

learning parameters are initialised at the same value and left to evolve as the number of spikes grows.

$$\text{P-STDP} = \Delta w_i = \begin{cases} \alpha_+ \cdot e^{-w_i}, & \leftrightarrow \delta \geq 0 & \text{LTP} \\ \alpha_-, & \leftrightarrow \delta < 0 & \text{LTD} \end{cases} \quad (1)$$

Where α_+ is the amplification parameter in the case of LTP and α_- is the amplification parameter in the case of LTD, their magnitudes being kept within a 4/3 ratio. The major input of P-STDP is its robustness to increased mathematical complexity in the neuron model [32] since performances can be preserved across several neuron models. P-STDP displayed robust results when shifting from a non-leaky IF to an Izhikevich neuron without hindering performances [32].

D. Reinforcement-STDP (R-STDP)

In biological learning, the brain's reward system is paramount in decision-making and neuromodulators such as dopamine modify synaptic behaviour [20]. Introduced by Mozafari et al. [20], this approach schematically relies on Pavlov's conditioning approach to learning which has been shown to contribute, together with Hebb's, to the neurobiological approach to learning and is nowadays widely used in Reinforcement Learning [29]. Namely from this perspective, spiking of \mathcal{N}_{aff} should be considered in a wider temporal span and not be limited to closely related spike times [10]. In other words, spikes having participated at some earlier point in the efferent neuron's firing should be entitled to some *reward* since there is a possibility that their action influenced the spike, this is called *eligibility trace* [31]. Here causality is tackled in terms of probability of participation. The main asset of R-STDP is the network learning *discriminating* features rather than *repeating* ones [20]. While classical STDP shines at differentiating highly distinct objects, performances can dwindle when dealing with objects presenting highly similar features. By contrast, R-STDP allows error rates to drop since the network is prompted to focus on diagnostic features which are not present in all images rather than only on those repeating [20]. P-STDP is used in combination with a WTA competition rule, the first neuron to spike being therefore also the only one to update its synaptic weights and the one which will impact the final decision of the network. Moreover, dropout can be used to force a maximal number of neurons into active classification. The learning rule uses α parameter to quantify the magnitude of weight change according to correlation and reward/punishment and η parameter as the adjustment factor to avoid over-fitting caused by unbalance in reward and punishment in training. The reinforcement award/punishment signal is given by last layer which compares the network's decision with its face value and sends the corresponding conditioning signal. There are four cases for STDP summed up in Table II, each formula corresponding to the following pattern with chosen α .

$$\eta_+ = \frac{|\text{missed samples}|}{|\text{training samples}|} \quad (2)$$

$$\eta_- = \frac{|\text{correctly labelled samples}|}{|\text{training samples}|} \quad (3)$$

	Reward signal	Punishment signal
Correlation ($\delta \geq 0$)	$\alpha_{r+} > 0$ η_+	$\alpha_{p+} > 0$ η_-
No correlation ($\delta < 0$) or silent pre-synaptic neuron	$\alpha_{r-} < 0$ η_+	$\alpha_{p-} < 0$ η_-

TABLE II
 α AND η R-STDP PARAMETERS ACCORDING TO CASE

$$\text{R-STDP} = \Delta w_i = \alpha \cdot w_i \cdot (1 - w_i) + \eta \cdot (\alpha \cdot w_i \cdot (1 - w_i)) \quad (4)$$

Note that ($|\text{missed}| + |\text{correctly labelled}|$) $\leq |\text{training}|$ since some samples might go un-handled.

E. Reversed-STDP (Rev-STDP)

In the biological world, Rev-STDP is a particularly complex type of STDP and though an extended demonstration falls outside the scope of this paper and can be found in Burbank's paper [6], we shall focus on the key points for its adaptation to computerised STDP. Rev-STDP occurs in top-down synapses in the brain, i.e. those synapses where communication happens in the opposite direction as is usual. This top-down communication often coexists with the more classical bottom-up one and therefore mechanisms of weight adaptation should take into account both the classical feed-forward communication and the feed-backward one. Some confusion might arise from the pre- and post- synaptic words, conceived as temporal references when they actually rely on topological description. In the brain topology, pre-synaptic –afferent– neurons are situated topologically *before* post-synaptic –efferent– ones. Therefore, when looking at spike timing, \mathcal{N}_{aff} should fire temporally before \mathcal{N}_{eff} for LTP to take place because that is the only case where there is a possibility for –though no insurance of– causality and not solely correlation. This happens so in the classical feed-forward journey of information. However, in the rarer feed-backward case, communication is carried out in the other way around. \mathcal{N}_{eff} –topologically situated *after* \mathcal{N}_{aff} – fires first at \mathcal{T}_0 therefore sending a message to the neuron that is *next* in its backward-going path: the \mathcal{N}_{aff} which receives the spike at \mathcal{T}_{+1} . This feed-backward communication often coexists with the classical feed-forward one and its modelling seems to result in more stable weight distributions when the learning is biased towards depression [6].

F. Stable-STDP (S-STDP)

A novel *homeostasis* parameter was introduced, which acted as an excitability indicator [22] and has a padding effect to adapt neural response to highly varying inputs. A post-synaptic neuron linked to highly active pre-synaptic neurons would have a low excitability since it would need more integration before firing, conversely a post-synaptic neuron linked to idle pre-synaptic ones would have a higher excitability in order to get the chance to fire at some point. This mechanism aims at handling the issue of slow and fast motion encoding of

AER. In this setting, all synaptic weights were initialised at the same constant, prior to the learning process. LTP and LTD were modified dynamically throughout the learning process in correlation with the variations of synaptic weights. As the synaptic weight grew larger, the effect of LTP decreased and LTD increased and conversely as synaptic weights diminished, the effect of LTD increased while that of LTP decreased, resulting in a smooth unimodal distribution. This particular process allows self-regulation of synaptic weights without implementing explicit bounds. S-STDP also relied on WTA rule maintained after the learning process was over. Similarly, the neuron model consisted of a refractory period during which afferent spikes had no effect.

G. Triplet-STDP (T-STDP)

This STDP is based on the main idea that motivation behind LTP and LTD lies not only in the relationship between a pair of post and pre synaptic spikes but between a triplet, either of two pre and one post or the other way around, though literature [28] suggests only the latter should be considered. Motivation for this rule is based on studying the evolution of potentiation with varying frequencies. Thinking in terms of sparse time windows, then an afferent spike’s *belonging* is unequivocal: it solely *belongs* to the nearest efferent neuron and should influence its potential and no other neuron’s potential. Yet complications arise when higher frequencies come into play, the more spikes there are, the harder it becomes to model which afferent spike should influence which efferent neuron and an LTP-favourable pre-post pair could eventually result in a virtual LTD-favourable post-pre pair [28] depending on the time window. Experiments have shown that pre-synaptic spikes which had participated in LTP for a particular efferent neuron and should have resulted in virtual LTD for other efferent neurons did not necessarily do so [28], hence suggesting that some mechanism allowed the brain to access the information of whether N_{eff} ’s potential had been increased or decreased: the triplet rule. In terms of implementation, this would suggest first, working with synaptic traces, second that two of those traces should be modelled. In his article, Krunglevicius showed that T-STDP helped achieve synaptic weight stability in the case of Schottky noise input [17].

III. PATTERN RECOGNITION TASKS

While many STDP benchmarks problems have been taken from the static image field therefore following in the footsteps of deep-learning, the recent interest in video input yielded better results and has become the focus of increasing attention in the field of SNNs. This being due to the temporal nature of STDP, making it a first class approach for sequential data rather than static.

Most static images are part of benchmark data sets such as Caltech or MNIST therefore presenting two main issues when compared with real-world images. First, the focus is often set on the target object, making it a central item while it is seldom the case in real-world images. Second, most images from these datasets are highly sanitised, therefore

leaving aside the question of separating details of changing backgrounds from the target object. Authors suggest that poor decreased performance in face vs background classification tasks could be caused by this phenomena [19]. Most of these datasets offer non real-life images containing only one main object abstracted from the real-world paraphernalia of details, which might have minor significance when thinking about transferability of neuron’s performance in detail-saturated real-world images such as ImageNet for instance.

A. Static grey-scale images

Masquelier and Thorpe [19] present in 2007 the motivation and main interest for combining STDP with temporal rank-order coding of information to promote unsupervised learning. Authors’ main focus was on feature extraction aiming to extract ten class-specific features (complex combination of edges) to which neurons belonging to the last layer before the classifier would be sensitive. Input stimuli was propagated in a discrete manner when real-life visual stream is continuous and using convolutions restricted area fields when real-life images would be wider, finally some feedback and top-down information probably made available by the cortex in terms of classification was overlooked. All in all, the extracted features allowed to differentiate and perform binary classification with results next to state of the art deep models, however these results should be considered thoroughly since the data-sets used consist in highly distinctive classes with almost no jitter nor noise. A major limitation was the absence of intra-class variance in the data-sets used, authors point out that classification of images presenting more variability in their essence (animals for instance) would require more training examples to be learnt since key features might not be repeated often enough in a smaller sample. Building invariance of shift and scale into the network increases space complexity, it sure decreases learning time but since one advantage of SNNs is to be energy efficient (as well as bio-like), it might be interesting to check the computational tradeoffs in both space and number of calculations. When comparing results obtained from interpreting binary detection, quantifying membrane potential or standard Hebbian approach, faces were consistently the best performing category and background the worst. We might infer that facial features are more distinguishable than others and also conclude that the background category was harder to detect because of its variability in content between images, which might not allow discriminating features to arise.

In the work by Mozafari et al. in 2018 [20], the learning paradigm presented is reinforcement, which aims at providing a full-SNN network without external classifier. The network presented is made of four layers A, B, C and D contrasted with a CNN with similar architecture but classic categorisation layer. Categorisation tasks with a focus on object invariance through perspective and lighting conditions. This article demonstrated that it is possible to implement full-SNN, locally-ruled networks without an external classifier if one resorts to RL. Moreover, robustness over visual variation has been demonstrated with ETH-80 and NORB data-sets.

Work conducted by Krunglevicius [17] in 2016 deals with the question of how small a spatial pattern could be relative to the total number of afferent neurons. When reducing input size, T-STDP outperformed other variants and could follow down to size 4 in a 64Hz/39Hz and size 8 in a 64Hz setting. However, diminishing the size of the input pattern resulted in loss of functionality since synaptic weights associated with pattern were insufficient to reach the threshold and produce a spike. Switches between LTP and LTD take place, which also happens in biological setting. The triplet rule hereby presented could be used to detect increase in firing rate instead of spatial pattern, in the case of very small patterns.

Multiple layer SNNs capable of processing large-scale static images were first introduced by Kheradpishesh et al. [15]. In their article, authors suggest that AER data could yield interesting results for further investigation using their multi-layered network. STDP learning took place in each convolutional layer, only starting in the $(n + 1)^{th}$ layer when it had finished in the n^{th} . WTA competition was used between neurons of the same convolutional layer. It was enforced globally for neurons pertaining to the same map and locally across maps. Since computation of each neuron’s variables is independent from its neighbours, convolution, pooling and STDP were performed in parallel on a GPU. Though DCNNs could outperform this network on large data-sets, they failed to do so in medium to small data-sets which suggests this network is more information-efficient since it manages to extract distinctive features earlier in the training.

Nessler et al. [21] discussed in 2009 the interpretation of synaptic weights as internal representations of salient visual features. This was also more recently pointed out by Tavanaei et al. [32] who showed that images could be reconstructed by convolving the final weight matrix. The simple STDP rule was sufficient since its more complex counterpart did not induce any significant improvement. The WTA rule implemented in the output layer induced over-specialisation of neurons which tried to be performant by reducing their scope. On the other hand, WTA favours online learning since neurons are encouraged to search for new items that they would be better fitted to detect. After learning, the internal weights represent a probability distribution and can therefore be used to compute a representation of the input data.

B. Static colour images

Static colour image recognition has been little exploited as of now in SNNs and pattern recognition in general. A detailed discussion of SNN vs auto-encoders approach can be found in Falez et al. [9]. Since including colour represented a relatively new challenge, authors experimented with several information combinations to see which performed best. Experiments showed that a combination of grey-scale and colour yielded higher accuracy rates, therefore this combination was adopted for the rest of the process. Temporal latency coding. Multiplicative STDP with layer lateral WTA inhibition was used with temporal latency information coding. While SNNs needed more features than auto-encoders, they did manage

reach performances within a 3σ of their auto-encoders counterparts with less features.

C. DVS camera and AER encoding

In their work, Bichler et al. [1] address two image-recognition tasks, both involving motion (trajectory detection) and based on AER data. Because the sensor relies on luminosity, it can model contrast. Robustness and tolerance are the main features put forth by the authors in this paper. Synaptic parameters could undergo a dispersion of 20% while neuronal parameters could undergo a dispersion of 10% and still be considered "good" 75% of the time. Bichler et al. also point that the size of the network was considerably smaller than classical networks for similar tasks. Limitations were particularly salient in the second task of car-counting since among the six lanes from which vehicles had to be counted, one was never learned by the network and at least one more lane was not learned 9% of the time. To account for this lack of learning, authors point at the role of random initialisation of synaptic weights.

Paredes-Vallès et al. [22] seek in 2019 to stabilise STDP through forcing synaptic weights to obey a unimodal distribution instead of the severely bounds-skewed one obtained with classical STDP where all synaptic weights either converge to 0 or 1, making it almost binary. The particularity of optical flow processing is that several updates can take place simultaneously in different locations and need to be processed. Their experiments used optical flow data in the form of event sequences, both synthetic (simulation of vertical and horizontal motion) and real (rotating disk and roadmap). Authors relied on the concept of synaptic trace and its stabilising role for synaptic weights and showed this resulted in increased synaptic equilibrium. Single-Synaptic Convolutional layer performed feature extraction while neighbourhood-WTA mechanism was used during learning but only neuron-specific WTA was maintained after. The synthetic data trial showed that out of 16 kernels, about half of them specialised in horizontal motion while the other half specialised in vertical motion and each direction of motion was captured by at least four kernels which specialised in different speed, therefore behaving like local velocity detectors.

IV. CHALLENGES OF BIO-INSPIRED PATTERN RECOGNITION

A. Biological plausibility

1) *Correlation vs Causality*: Among the most famous debates in statistical ML is the correlation vs causality question, here focused on the value of δ . Reinforcing synapses is based on detecting correlations [15] and seems to adopt a classical approach where any chronological hierarchy is banned. This makes sense in purely mathematical modelling of STDP. However, behind detecting correlation, what is intended is often in fact to detect the causality rule lying behind it, correlation only being the tangible trace of its more elusive cousin. Despite clear assertion that correlation does not imply causation [4], [13], the latter still represents the *holy grail* of learning since

finding causal features would ensure better prediction. Because causal inference is highly complex and only partially observable in ML settings, research focuses on tracking associational inference. Causality requires ruling out other causes and being able to clearly differentiate \mathcal{A} from $\neg\mathcal{A}$, which implies a deeper understanding than the one accessible in current tasks. Withal statistical causal inference relying on experimental practices might be an interesting solution to this ordeal [13]. Though it might be considered sufficient to rely on correlation for learning, causality would be a sturdier alternative if correctly determined. Despite often settling for correlation, most algorithms aspire to detect causation. In the biological world and assuming time as linear, an event \mathcal{E}_i taking place *before* another event \mathcal{E}_{i+1} is a possible candidate for causality with $\mathcal{E}_i \Rightarrow \mathcal{E}_{i+1}$, contrarily \mathcal{E}_{i+1} can not be a causal candidate hence $\mathcal{E}_i \Leftarrow \mathcal{E}_{i+1}$ is false. A positive δ corresponds to \mathcal{E}_i taking place prior to \mathcal{E}_{i+1} on the linear time continuum, therefore allowing for a possible causality on top of correlation. On the other hand a negative δ corresponds to \mathcal{E}_i taking place after \mathcal{E}_{i+1} which prevents it from being a causal candidate while still allowing for association. Therefore, while $\delta \geq 0$ does not guarantee causation, $\delta < 0$ guarantees its absence, ergo if one tries to aim for biological plausibility, approaches relying on $\delta < 0$ should be discarded.

2) *Locality*: Simplifications required by computer simulations make most biological phenomena overly complex. Synaptic processes in the brain occur locally in general, hence any biologically plausible learning rule must be bound by some topological limitation. This makes multiple-layered SNNs particularly difficult to implement. Moreover, certain learning rules such as reversed or mirror STDP are sometimes implemented in a way which does not fully coincide with the current knowledge of neurobiology. Synaptic weight sharing, though useful, since it allows to ensure location invariance in pattern recognition [22], is unlikely to take place in the brain [19] since it is a non local process and another approach relying only on local weight-sharing would allow for increased biological plausibility.

B. Generalising potential

Invariance is the key to generalisation. Focus is set on several aspects of invariance, such as scale and shift [1] or perspective and lighting conditions [20]. One way to tackle invariance of shift and scale is through duplication of cells at all positions and scales, therefore building in the structure of the network rather than relying on training. This allows the number of training examples to be reduced yet can not be considered as biologically plausible [1]. Size invariance can represent a major hurdle since in some cases hyper-parameters required optimisation for each pattern size [17], implying that the network was unable to generalise through size. The question of bias in image data-sets has been abundantly addressed, be it referring to social [16], [35] or geometrical bias [33] such as object-centred point of view, which are common in staple datasets used for benchmark tasks in pattern recognition. The network’s ability to generalise features learnt

is often at stake. Categorisation tasks such as Cal-tech’s face-motorbike focus on salient and common features [15] ignoring the drastically changing backgrounds. Taking into account the object-centric, canonical angle [33] of these data-sets, the evolution of such results when faced with a wider variety of angles and motorbike models all taken with the exact same background would be of interest. Similarly, discrimination tasks such as gender identification can represent hurdles for SNNs which rely heavily on common features [15], [20] rather than on diagnostic ones [20].

C. Information processing and coding

Image processing does not directly deal with real data but with a simplified representation of it. This is as true in artificial setting as it is in biological. That being said, the goal is to encode reality in a way that allows for maximal representativeness, minimal space complexity and minimal loss. The two widely used options are rate and temporal coding while population coding remains liminal, we refer the reader to Brette’s paper [4] for an extensive discussion on this topic. Issues arise right from encoding where the ON-centre/OFF-centre image processing, for instance, results in information loss [9]. Widely used ON-centre/OFF-centre and DoG filters [9], [15] rely on edges to encode input spikes therefore biasing filters towards edge-dominated features [9]. This, while allowing for enhanced performance in shape detection, puts SNNs at a disadvantage when tackling colour input [9] In the same way, classical SNN pattern recognition only relies on grey-scale images, therefore losing all colour information [9]. Moreover, colour is not the only factor that should be taken into account since dimensions of input pattern must meet some criteria. Size should remain reasonable since diminishing it too drastically ultimately results in neurons failing to act selectively for combined synaptic signals’ strength associated with pattern are insufficient to cause a spike [17]. In terms of coding, differentiating horizontal and vertical edges from other edges could prove optimal since both of these types are common and seldom incarnate diagnostic features [19]. The latter is coherent with the integration of spikes equations, since \mathcal{N}_{eff} ’s threshold shan’t be lowered to avoid spiking because of noise, too few \mathcal{N}_{aff} spiking with their favourite pattern can not trigger a spike. Finally, AER encoding though better suited to SNN’s temporal abilities, represents a specific challenge since it is based on events. Data encoding therefore depends on both the firing rate of the sensor and on the optical flow and sensitivity settings of the camera. Fast and slow motion encoding imply that the rate of input data can vary, hence calling for adapted ω mechanisms [22].

CONCLUSION

This paper provides the first STDP-centred review in spiking networks-based pattern recognition, to the best of our knowledge. A synthesis is given in Table I. SNN and STDP benefit from a range of promising features for pattern recognition and yet, a number of challenges lie ahead before they become a realistic alternative to deep CNN. The effective use of SNNs to

tackle modern pattern recognition problems is promising and yet still in its infancy. We suggest the choice of the STDP rule bears significant impact in the networks' performance. Future work focusing on the interplay between the type of neuron and the choice of the STDP would bring decisive insight.

Though the most widely used type of STDP in SNNs for computer-vision has been the C-STDP and particularly its multiplicative variant which allows pondering, we advocate that marginally explored STDP variants such as the probabilistic [32] or reinforced one [20] deserve increased attention. R-STDP presented by Mozafari et al. [20] yields promising results, in particular when dealing with the extraction of discriminative features. This STDP has received attention in fields related to control theory such as autonomous vehicles [2] and robotics [3], [27] where it allows for the advantages of reinforcement learning to be combined with bio-inspiration. As for Tavanaei et al.'s [32] P-STDP, the demonstrated robustness against complex bio-inspired neurons is a valuable advantage for future implementations willing to rely on complex neuron models. An implementation of an activity-based P-STDP has been exploited in reservoir computing [14] and displayed the best performance, irrelevant of reservoir size.

Because of its transferability to other fields and combination with the strong versatility and performances reinforcement learning offers, we believe R-STDP could well become the centre of attention in coming years and probably supplant its current competitor C-STDP.

REFERENCES

- [1] O. Bichler, D. Querlioz, S. J. Thorpe, J. P. Bourgoin, and C. Gamrat. Extraction of temporally correlated features from dynamic vision sensors with spike-timing-dependent plasticity. *Neural Networks*, 32:339–348, aug 2012.
- [2] Z. Bing, C. Meschede, K. Huang, G. Chen, F. Rohrbain, M. Akl, and A. Knoll. End to end learning of spiking neural network based on r-stdp for a lane keeping vehicle. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4725–4732, May 2018.
- [3] Z. Bing, C. Meschede, F. Röhrebein, K. Huang, and A. C. Knoll. A survey of robotics control based on learning-inspired spiking neural networks. *Frontiers in neurobotics*, 12:35–35, Jul 2018. 30034334[pmid].
- [4] R. Brette. Philosophy of the spike: Rate-based vs. spike-based theories of the brain. *Frontiers in Systems Neuroscience*, 9:151, 2015.
- [5] K. S. Burbank. Mirrored stdp implements autoencoder learning in a network of spiking neurons. *PLoS computational biology*, 11(12):e1004566–e1004566, Dec 2015. 26633645[pmid].
- [6] K. S. Burbank and G. Kreiman. Depression-biased reverse plasticity rule is required for stable learning at top-down connections. *PLOS Computational Biology*, 8(3):1–16, 03 2012.
- [7] N. Caporale and Y. Dan. Spike timing-dependent plasticity: A hebbian learning rule. *Annual Review of Neuroscience*, 31(1):25–46, 2008.
- [8] F. Danneville, C. Loyez, K. Carpentier, I. Sourikopoulos, E. Mercier, and A. Cappy. A Sub-35 pW Axon-Hillock artificial neuron circuit. *Solid-State Electronics*, 153:88–92, mar 2019.
- [9] P. Falez, P. Tirilly, I. M. Bilasco, P. Devienne, and P. Boulet. Unsupervised visual feature learning with spike-timing-dependent plasticity: How far are we from traditional feature learning approaches? *Pattern Recognition*, 93:418–429, sep 2019.
- [10] A. Harry Klopff. A neuronal model of classical conditioning. *Psychobiology*, 16(2):85–125, Jun 1988.
- [11] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR 2016*, 2016.
- [12] D. O. Hebb. The organization of behaviour, a neuropsychological theory. *The American Journal of Psychology*, 63(4):633–642, 1950.
- [13] P. W. Holland. Statistics and causal inference. *Journal of the American Statistical Association*, 81(396):945–960, 1986.
- [14] Y. Jin and P. Li. Ap-stdp: A novel self-organizing mechanism for efficient reservoir computing. In *2016 International Joint Conference on Neural Networks (IJCNN)*, pages 1158–1165, July 2016.
- [15] S. R. Kheradpisheh, M. Ganjtabesh, S. J. Thorpe, and T. Masquelier. STDP-based spiking deep convolutional neural networks for object recognition. *Neural Networks*, 99:56–67, mar 2018.
- [16] A. Khosla, T. Zhou, T. Malisiewicz, A. A. Efros, and A. Torralba. Undoing the damage of dataset bias. In A. Fitzgibbon, S. Lazebnik, P. Perona, Y. Sato, and C. Schmid, editors, *Computer Vision – ECCV 2012*, pages 158–171, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.
- [17] D. Krunglevicius. Modified STDP Triplet Rule Significantly Increases Neuron Training Stability in the Learning of Spatial Patterns. *Advances in Artificial Neural Systems*, 2016:1–12, aug 2016.
- [18] W. Maass. Networks of spiking neurons: The third generation of neural network models. *Neural Networks*, 10(9):1659–1671, 1997.
- [19] T. E. Masquelier and S. J. Thorpe. Unsupervised Learning of Visual Features through Spike Timing Dependent Plasticity. *PLOS Computational Biology*, 2007.
- [20] Mozafari, Kheradpisheh, Masquelier, Nowzari-Dalini, and Gantabesh. First-spike-based visual categorization using reward-modulated STDP. *IEEE Transactions on Neural Networks and Learning Systems*, 29(12):6178–6190, dec 2018.
- [21] B. Nessler, M. Pfeiffer, and W. Maass. Stdp enables spiking neurons to detect hidden causes of their inputs. In *Advances in Neural Information Processing Systems 22*.
- [22] F. Paredes-Vallés, K. Y. W. Scheper, and G. C. H. E. de Croon. Unsupervised Learning of a Hierarchical Spiking Neural Network for Optical Flow Estimation: From Events to Global Motion Perception. *Pattern Analysis and Machine Intelligence*, jul 2018.
- [23] J. Pei, L. Deng, S. Song, and M. Z. et al. Towards artificial general intelligence with hybrid Tianjic chip architecture. *Nature*, 572, 2019.
- [24] E. T. Rolls and T. Milward. A model of invariant object recognition in the visual system: Learning rules, activation functions, lateral inhibition, and information-based performance measures. *Neural Computation*, 12(11):2547–2572, 2000.
- [25] K. Roy, A. Jaiswal, and P. Panda. Towards spike-based machine intelligence with neuromorphic computing. *Nature*, 575, 2019.
- [26] A. Sboev, D. Vlasov, A. Serenko, R. Rybka, and I. Moloshnikov. A comparison of learning abilities of spiking networks with different spike timing-dependent plasticity forms. *Journal of Physics: Conference Series*, 681:012013, feb 2016.
- [27] M. S. Shim and P. Li. Biologically inspired reinforcement learning for mobile robot collision avoidance. In *2017 International Joint Conference on Neural Networks (IJCNN)*, pages 3098–3105, May 2017.
- [28] J. Sjöström and W. Gerstner. Spike-timing dependent plasticity. *Scholarpedia*, 5(2):1362, 2010. revision #184913.
- [29] R. S. Sutton and A. G. Barto. *Reinforcement Learning, An Introduction*. 2018.
- [30] A. Taherkhani, A. Belatreche, Y. Li, G. Cosma, L. P. Maguire, and T. M. McGinnity. A review of learning in biologically plausible spiking neural networks. *Neural Networks*, 12, 2020.
- [31] B. Tanner and R. S. Sutton. Td(λ) networks: Temporal-difference networks with eligibility traces. In *Proceedings of the 22nd International Conference on Machine Learning, ICML '05*, page 888–895, New York, NY, USA, 2005. Association for Computing Machinery.
- [32] A. Tavanaei, T. Masquelier, and A. S. Maida. Acquisition of visual features through probabilistic spike-timing-dependent plasticity. In *Proceedings of the International Joint Conference on Neural Networks*, volume 2016-October, pages 307–314. Institute of Electrical and Electronics Engineers Inc., oct 2016.
- [33] A. Torralba and A. A. Efros. Unbiased look at dataset bias. In *CVPR 2011*, pages 1521–1528, June 2011.
- [34] S. J. Verzi, F. Rothganger, O. D. Parekh, and T.-T. Quach. Computing with spikes: The advantage of fine-grained timing. *Neural Computation*, 30(10), 2018.
- [35] T. Wang, J. Zhao, M. Yatskar, K.-W. Chang, and V. Ordonez. Balanced datasets are not enough: Estimating and mitigating gender bias in deep image representations. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.
- [36] X. Xie, R. H. R. Hahnloser, and H. S. Seung. Learning winner-take-all competition between groups of neurons in lateral inhibitory networks. In T. K. Leen, T. G. Dietterich, and V. Tresp, editors, *Advances in Neural Information Processing Systems 13*, pages 350–356. MIT Press, 2001.