



## OPEN ACCESS

## EDITED BY

Mostafa Rahimi Azghadi,  
James Cook University, Australia

## REVIEWED BY

Ben Walters,  
James Cook University, Australia  
Anthony Maida,  
University of Louisiana at Lafayette,  
United States

## \*CORRESPONDENCE

Ioan Marius Bilasco  
✉ marius.bilasco@univ-lille.fr

RECEIVED 15 March 2024

ACCEPTED 11 July 2024

PUBLISHED 24 July 2024

## CITATION

Goupy G, Tirilly P and Bilasco IM (2024) Paired competing neurons improving STDP supervised local learning in Spiking Neural Networks. *Front. Neurosci.* 18:1401690. doi: 10.3389/fnins.2024.1401690

## COPYRIGHT

© 2024 Goupy, Tirilly and Bilasco. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/). 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.

# Paired competing neurons improving STDP supervised local learning in Spiking Neural Networks

Gaspard Goupy, Pierre Tirilly and Ioan Marius Bilasco\*

Univ. Lille, CNRS, Centrale Lille, UMR 9189 CRISTAL, Lille, France

Direct training of Spiking Neural Networks (SNNs) on neuromorphic hardware has the potential to significantly reduce the energy consumption of artificial neural network training. SNNs trained with Spike Timing-Dependent Plasticity (STDP) benefit from gradient-free and unsupervised local learning, which can be easily implemented on ultra-low-power neuromorphic hardware. However, classification tasks cannot be performed solely with unsupervised STDP. In this paper, we propose Stabilized Supervised STDP (S2-STDP), a supervised STDP learning rule to train the classification layer of an SNN equipped with unsupervised STDP for feature extraction. S2-STDP integrates error-modulated weight updates that align neuron spikes with desired timestamps derived from the average firing time within the layer. Then, we introduce a training architecture called Paired Competing Neurons (PCN) to further enhance the learning capabilities of our classification layer trained with S2-STDP. PCN associates each class with paired neurons and encourages neuron specialization toward target or non-target samples through intra-class competition. We evaluate our methods on image recognition datasets, including MNIST, Fashion-MNIST, and CIFAR-10. Results show that our methods outperform state-of-the-art supervised STDP learning rules, for comparable architectures and numbers of neurons. Further analysis demonstrates that the use of PCN enhances the performance of S2-STDP, regardless of the hyperparameter set and without introducing any additional hyperparameters.

## KEYWORDS

Spiking Neural Networks, image recognition, supervised STDP, Winner-Takes-All, intra-class competitive learning

## 1 Introduction

Artificial Neural Networks (ANNs) have gathered exponential attention across diverse domains in recent years (Abiodun et al., 2018). However, ANN training suffers from high and inefficient energy consumption on modern computers based on the von Neumann architecture (Zou et al., 2021). Spiking Neural Networks (SNNs) (Ponulak and Kasinski, 2011) implemented on neuromorphic hardware (Schuman et al., 2017; Shrestha et al., 2022) have emerged as a promising solution to overcome the von Neumann bottleneck (Zou et al., 2021) and enable energy-efficient computing. In particular, memristive-based neuromorphic hardware (Jeong et al., 2016; Xu et al., 2021) is an excellent candidate for ultra-low-power applications, potentially reducing energy consumption by at least one order of magnitude compared to state-of-the-art CMOS-based neuromorphic hardware (Milo et al., 2020; Liu et al., 2021a), and by several orders of

magnitude compared to GPUs (Yao et al., 2020; Li et al., 2022). However, direct training of SNNs on neuromorphic hardware comes with a major constraint: implementing network-level communication is challenging and requires significant circuitry overhead (Zenke and Nefcici, 2021). Therefore, the involved learning mechanisms should rely on local weight updates, i.e., updates based only on the two neurons that the synapse connects.

Training SNNs to achieve state-of-the-art performance is typically accomplished with adaptations of Backpropagation (BP) (Eshraghian et al., 2021; Dampfhofer et al., 2023). However, these methods are challenging to implement on neuromorphic hardware as they rely on non-local learning. In addition, they employ gradient approximation to circumvent the non-differentiable nature of the spike generation function, which is suboptimal. Other approaches attempted to make gradient computation local, notably by utilizing feedback connections (Nefcici et al., 2017; Zenke and Ganguli, 2018), or by employing a layer-wise cost function (Ma et al., 2021; Mirsadeghi et al., 2021). Yet, they do not solve the gradient approximation problem. Furthermore, all of the aforementioned BP-based methods rely solely on supervised learning, which increases the dependence on labeled data. We believe that machine learning algorithms should minimize this dependence on supervision by employing unsupervised feature learning (Bengio et al., 2013). Hence, an optimal classification system should include both unsupervised and supervised components, for data representation and classification, respectively.

Spike Timing-Dependent Plasticity (STDP) (Caporale and Dan, 2008) is a gradient-free, unsupervised and local alternative to BP, inspired by the principal form of plasticity observed in biological synapses (Hebb, 1949). STDP solves the previously mentioned limitations of BP and is inherently implemented in memristor circuits (Querlioz et al., 2011; Schuman et al., 2017), which makes it suitable for on-chip training on memristive-based neuromorphic hardware (Saighi et al., 2015; Khacef et al., 2023). Unsupervised feature learning with STDP has been extensively studied in the literature, particularly for image recognition tasks. Convolutional SNNs (CSNNs) trained with STDP have demonstrated the ability to improve data representation by extracting relevant features from images (Tavanaei and Maida, 2017; Ferré et al., 2018; Kheradpisheh et al., 2018; Falez et al., 2019b; Srinivasan and Roy, 2019). However, to perform classification based on the extracted features, these solutions employ external classifiers, such as ANNs or support vector machines, which are incompatible with neuromorphic hardware. To leverage the potential of these CSNNs in neuromorphic hardware and enable end-to-end SNN solutions, spiking classifiers trained with local supervised learning rules must be designed. Ensuring compatibility between classifiers and CSNNs, particularly regarding the type of local rule employed, could significantly mitigate hardware implementation overhead.

Although STDP is traditionally formulated for unsupervised learning, it can be adapted for supervised learning by incorporating a third factor, taking the form of an error signal that is used to guide the STDP updates (Frémaux and Gerstner, 2015). As a result, STDP enables end-to-end SNNs to perform classification tasks by combining unsupervised STDP for feature extraction and

supervised STDP for classification (Shrestha et al., 2017; Thiele et al., 2018; Lee et al., 2019; Mozafari et al., 2019). Several supervised adaptations of STDP are reported in the literature (Ponulak and Kasiński, 2010; Shrestha et al., 2017, 2019; Lee et al., 2019; Tavanaei and Maida, 2019; Hao et al., 2020; Zhao et al., 2020; Zhang et al., 2021; Saranirad et al., 2022). Yet, all of the aforementioned rules are designed to train SNNs with multiple spikes per neuron, which is undesirable because state-of-the-art CSNNs trained with unsupervised STDP usually employ one spike per neuron. For compatibility with these CSNNs, a spiking classifier trained with supervised STDP should adhere to this single-spike approach. In addition, it has been shown that using one spike per neuron with temporal coding presents several advantages for visual tasks, including fast information transfer, low computational cost, and improved energy efficiency (Rullen and Thorpe, 2001; Park et al., 2020; Guo et al., 2021). The literature exploring supervised adaptations of STDP for training SNNs with only one spike per neuron is limited. Reward-modulated STDP (R-STDP) (Mozafari et al., 2019) is a learning rule based on Winner-Takes-All (WTA) competition (Ferré et al., 2018) that modulates the polarity of the STDP update to apply a reward or a punishment. R-STDP has gained popularity notably for its simplicity, but it results in inaccurate weight updates as only the polarity of STDP is adjusted. Recently, Supervised STDP (SSTDP) (Liu et al., 2021b) proposes a method to modulate, in the output layer, both the polarity and intensity of STDP with temporal errors, resulting in more accurate weight updates. When combined with the non-local optimization process of BP, SSTDP enables state-of-the-art performance in deep SNNs on various image recognition datasets. However, it has not yet been investigated in settings based on local learning, combining unsupervised STDP for feature extraction and SSTDP for classification. In addition, we claim that SSTDP faces two issues that may limit its performance. First, SSTDP training results in a limited number of STDP updates per epoch, which can lead to premature training convergence. Second, SSTDP training causes the saturation of firing timestamps toward the maximum firing time, which can limit the ability of the SNN to separate classes.

In this paper, we focus on the supervised STDP training of a spiking classification layer with one spike per neuron and temporal decision-making. This classification layer is the output layer of an SNN equipped with unsupervised STDP for feature extraction, as illustrated in Figure 1. The main contributions of this paper include the following:

1. In a preliminary study, we analyze the behavior of SSTDP when used to train the classification layer of an SNN equipped with unsupervised STDP. We demonstrate that the rule encounters two issues that may limit its performance: the limited number of STDP updates per epoch and the saturation of firing timestamps toward the maximum firing time.
2. To address the issues of SSTDP, we propose Stabilized Supervised STDP (S2-STDP), a supervised STDP learning rule that teaches neurons to align their spikes with dynamically computed desired timestamps derived from the average firing time within the layer.
3. To further enhance the learning capabilities of our classification layer trained with S2-STDP, we introduce a training architecture

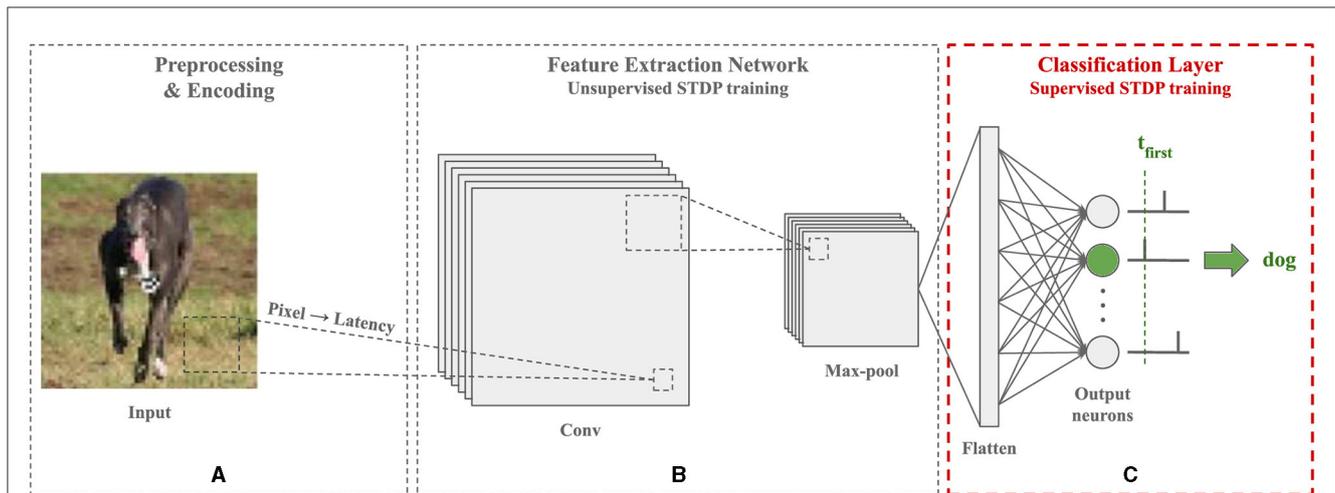


FIGURE 1

Architecture of the SNN employed in this paper for image recognition tasks. First, the image is preprocessed and each pixel is encoded into a single floating-point timestamp using the latency coding scheme. Then, a Convolutional SNN (CSNN) trained with unsupervised STDP is used to extract relevant features from the image. The resulting feature maps are compressed through a max-pooling layer to reduce their size and provide invariance to translation on the input. Lastly, they are flattened and fed to a fully-connected SNN trained with a supervised adaptation of STDP for classification. Each output neuron is associated with a class and the first one to fire predicts the label. Training is done in a layer-wise fashion. This classification pipeline organized into three blocks provides a flexible framework for SNNs combining feature extraction and classification. In this paper, we focus on the classification layer block (C), which may be integrated after other encoding or feature extraction blocks based on latency coding.

called Paired Competing Neurons (PCN). This method associates each class with paired neurons and encourages neuron specialization toward target or non-target samples through intra-class competition.

- We evaluate the performance of S2-STDP and PCN on three image recognition datasets of growing complexity: MNIST, Fashion-MNIST, and CIFAR-10.

The remainder of this paper is organized as follows. In Section 2, we provide the necessary background information about the SNN employed in this study. In Section 3, we demonstrate experimentally the aforementioned issues of SSTDP, which we address with our contributions. In Section 4, we describe our spiking classification layer and our proposed training methods. In Section 5, we cover our results on image recognition datasets and provide an in-depth investigation of the key characteristics of our methods. In Section 6, we conclude the paper. The source code is publicly available at: <https://gitlab.univ-lille.fr/fox/snn-pcn>.

## 2 Background

### 2.1 Neural coding

Since spiking neurons communicate through spikes, encoding the image, as illustrated in Figure 1A, is a necessary step before the SNN can process it. In this work, we use a temporal coding scheme called latency coding (Thorpe et al., 2001). This scheme represents each pixel by a single spike timestamp, thus limiting the number of generated spikes and making the coding more energy efficient, which is suitable for implementation on ultra-low-power devices. For a given pixel  $x \in [0, 1]$ , we calculate its spike timestamp  $t(x)$  as

following Equation 1:

$$t(x) = T_{\max} \cdot (1 - x), \quad (1)$$

where  $T_{\max}$  is the maximum firing time (set to 1 in this work). Consequently, the intensity of the pixel is encoded through latency: higher pixel values correspond to lower latencies, and vice versa. Spike timestamps are represented by floating-point values to align with spike-based communication in event-driven neuromorphic hardware.

### 2.2 Neuron model

To simulate the dynamics of spiking neurons, we use the Single-Spike Integrate-and-Fire (SSIF) model (Kheradpisheh and Masquelier, 2020; Goupy et al., 2023), describing IF neurons that can fire at most once per sample. Following latency coding, SSIF neurons encode the magnitude of their activation through spike timing: neurons firing first are the most strongly activated. The neurons integrate input spikes to their membrane potential  $V$  with Equation 2:

$$\frac{\partial V_j(t)}{\partial t} = \sum_i w_{ij} \cdot S_i(t), \quad (2)$$

where  $t$  is the timestamp,  $V_j$  is the membrane potential of neuron  $n_j$ ,  $w_{ij}$  is the weight of the synapse from  $n_i$  to  $n_j$ , and  $S_i(t)$  indicates the presence ( $S_i(t) = 1$ ) or absence ( $S_i(t) = 0$ ) of a spike from input neuron  $n_i$  at timestamp  $t$ . When the membrane potential of neuron  $n_j$  exceeds a defined threshold  $V_{th}$  (i.e.  $V_j(t) \geq V_{th}$ ), the neuron emits a spike at timestamp  $t$  and is deactivated until the next sample is shown.

## 2.3 Unsupervised STDP feature learning

Unsupervised training of CSNNs with STDP, as illustrated in Figure 1B, is an effective approach for improving image representation before classification without using any labeled data. In this work, we use the CSNN model introduced by Falez et al. (2019b). This model comprises trainable convolutional layers that extract spatial features from an encoded image, and non-trainable max-pooling layers that reduce the size of the feature maps and provide translation invariance on the input. The output of the CSNN consists of two-dimensional feature maps containing the output spikes of the neurons in the final layer. Following the SSIF neuron model, there is at most one spike per neuron (i.e., per feature) and the intensity of the activation is encoded temporally.

STDP training consists in adjusting synaptic weights of the convolutional layers based on the time difference between input and output neuron spikes. The rule considers both causal and non-causal relationships, increasing the synaptic weight when the input neuron fires before the output neuron (causal), and decreasing it otherwise (non-causal). For an output neuron  $n_j$ , its weights are updated with the multiplicative STDP (Querlioz et al., 2011) following Equation 3:

$$\Delta w_{ij} = \begin{cases} A^+ \times \exp\left(-\beta \frac{w_{ij} - w_{\min}}{w_{\max} - w_{\min}}\right) & \text{if } t_j \geq t_i \\ A^- \times \exp\left(-\beta \frac{w_{\max} - w_{ij}}{w_{\max} - w_{\min}}\right) & \text{o.w.} \end{cases}, \quad (3)$$

where  $w_{ij}$  represents the weight of the synapse connecting input neuron  $n_i$  and output neuron  $n_j$ ,  $\Delta w_{ij}$  is the weight change,  $A^+$  and  $A^-$  are the positive and negative learning rates,  $\beta$  is the saturation factor,  $w_{\min}$  and  $w_{\max}$  are the minimum and maximum achievable weight values in the layer, and  $t_i$  is the firing timestamp of neuron  $n_i$ . In addition to STDP, WTA competition and homeostatic plasticity are employed in the convolutional layers (Falez et al., 2019b). Via lateral inhibition, WTA competition promotes the learning of various patterns: for each sample, only the weights of the first neuron to fire are updated. Via threshold adaptation, homeostatic plasticity regulates the WTA competition.

We recall that we focus, in this work, on the spiking classification layer, given the extensive research already existing on CSNNs trained with unsupervised STDP. For an in-depth explanation and investigation of the employed CSNN model, we refer readers to (Falez et al., 2019a,b, 2020). In addition, other feature extractors may be employed instead of this model, provided that the resulting output features adhere to the principles of the SSIF model.

## 2.4 Supervised STDP for classification

Supervised training of SNNs with STDP, as illustrated in Figure 1C, incorporates the class information of the input sample to adjust the STDP updates. Given the architecture of the SNN employed in this work, we focus on supervised STDP at the classification (i.e., output) layer only. The literature on supervised STDP training includes various approaches. Shrestha et al. (2017); Lee et al. (2019); Hao et al. (2020) use teacher neurons that push target neurons to fire and non-target neurons to remain silent.

Neurons that fire update their weights using unsupervised STDP, hence without error modulation and only on samples of their class. Shrestha et al. (2019); Zhao et al. (2020); Zhang et al. (2021) derive errors based on spike counts to modulate the sign and intensity of the STDP updates. However, in the context of temporal learning with one spike per neuron, this is equivalent to using ternary errors (+1, -1, 0), where only the sign of the update is adjusted. Tavanaei and Maida (2019) define a target spike train and update, following each target spike, the weights of all neurons with an STDP rule modulated by a ternary error. When using one spike per neuron, the target spike train consists of a single spike and produces a single sign-modulated update per neuron. Most of these rules are designed and effective for training SNNs with multiple spikes and weight updates per neuron. In this work, we employ temporal coding with one spike per neuron. Hence, the supervised STDP rule employed for training the classification layer must be effective with this single-spike constraint. The remainder of this section presents the main existing rules designed for training a spiking classification layer with one spike per neuron.

### 2.4.1 R-STDP

Reward-Modulated STDP (R-STDP) (Mozafari et al., 2019) is a simple learning rule that modulates the sign of the STDP update through rewards and punishments. R-STDP is combined with WTA competition such that, for each sample, only the first neuron to fire, denoted  $n_{j^*}$ , receives a weight update. The error  $e_{j^*}$  of  $n_{j^*}$  is given by Equation 4:

$$e_{j^*} = \begin{cases} +1 & \text{if } c_{j^*} = y \\ -1 & \text{if } c_{j^*} \neq y \end{cases}, \quad (4)$$

where  $c_{j^*}$  is the class of neuron  $n_{j^*}$ , and  $y$  is the class of the sample. In practice, R-STDP requires mapping each class to multiple neurons to achieve satisfactory performance. The rule is also employed with adaptive learning rates and dropout to reduce overfitting.

### 2.4.2 SSTDP

Supervised STDP (SSTDP) (Liu et al., 2021b) is a learning rule with state-of-the-art performance, modulating both the sign and intensity of the STDP update with temporal errors. SSTDP training consists in teaching output neurons (one per class) to fire within one of two desired time ranges. For each sample, these time ranges are dynamically computed for the target neuron (i.e. the neuron associated with the class of the sample) and the non-target neurons (i.e. the neurons not associated with the class of the sample), based on the average firing time  $T_{\text{mean}}$  in the layer. At the end of the presentation of a sample, each output neuron updates its weights with an error-modulated STDP. The error  $e_j$  of neuron  $n_j$  is defined as the temporal difference between the neuron firing timestamp and its time range boundary, as described by Equation 5:

$$e_j = \begin{cases} \max\{0, t_j - (T_{\text{mean}} - g_1)\} & \text{if } c_j = y \\ \min\{0, t_j - (T_{\text{mean}} + g_2)\} & \text{if } c_j \neq y \end{cases}, \quad (5)$$

where  $t_j$  is the firing timestamp of neuron  $n_j$ ,  $g_1$  and  $g_2$  are defined time gaps that control the distance from  $T_{\text{mean}}$ ,  $c_j$  is the class of

neuron  $n_j$ , and  $y$  is the class of the sample. In other words, the time range of the target neuron ( $c = y$ ) is  $[0, T_{\text{mean}} - g_1]$  and the time range of the non-target neurons ( $c \neq y$ ) is  $[T_{\text{mean}} + g_2, T_{\text{max}}]$ . If a neuron fires within its desired range, the error is zero and the neuron weights are not updated. Note that, due to the min and max functions, SSTDP exclusively permits positive errors for target neurons (promoting earlier firing) and negative errors for non-target neurons (promoting later firing), potentially restricting accurate control over the output firing timestamps.

### 3 Preliminary study

In Section 2.4.2, we briefly outlined the SSTDP training method for an output layer with one spike per neuron. While this rule achieves state-of-the-art performance, the formulation of the desired time ranges (as presented in Equation 5) raises a concern: SSTDP offers restricted control over the output firing timestamps. To demonstrate this, we conducted preliminary experiments on the MNIST (LeCun et al., 1998) and the Fashion-MNIST (Xiao et al., 2017) datasets, using SSTDP to train the classification layer of an SNN equipped with unsupervised STDP, as illustrated in Figure 1. It led us to identify two primary issues that we aim to resolve in this work:

1. the limited number of STDP updates per epoch;
2. the saturation<sup>1</sup> of firing timestamps toward the maximum firing time.

First, if a neuron fires within its desired time range during training, its weights are not updated. Since neurons can easily reach their desired time range, training with SSTDP results in a limited number of updates per epoch. Figure 2 illustrates the update ratio per epoch, computed as the average number of updates per neuron divided by the number of samples. For both datasets, the total update ratio is around 50% in the initial epoch, implying that, for each sample, only half of the neurons receive a weight update. As the number of epochs increases, this ratio quickly decreases to 9% for MNIST and 16% for Fashion-MNIST. This limited number of updates leads to rapid training convergence, as indicated by the stabilization of training accuracies within a few epochs. However, since many samples are not involved in the training process, we believe that such rapid convergence is premature and may reduce the capabilities of the SNN to generalize, resulting in suboptimal model performance. In addition, this training process is inefficient because many samples undergo computational processing by the SNN without producing weight updates.

Second, because non-target updates are more frequent than target updates (see Figure 2), neurons are continually pushed to fire later. It creates a saturation effect where the firing timestamps of neurons, as training continues, progressively approach the maximum firing time. As observed in Figure 3, the average firing time grows rapidly in the first epochs and then stabilizes close to the maximum firing time (set to 1). During the last training epoch, we compute an average firing time of  $0.98 \pm 0.009$  for MNIST

and  $0.99 \pm 0.008$  for Fashion-MNIST. Note that the standard deviation corresponds to the variability in the average firing time between samples: it indicates how much the average firing time varies across samples during an epoch. We believe that the observed saturation effect limits the expressivity of the SNN (i.e., its ability to capture discriminant spatiotemporal information) and defeats the principles of temporal coding as the majority of the input spikes are integrated by the output neurons. In addition, the low variance in the average firing time across samples of different classes may affect the ability of the SNN to separate classes.

### 4 Methods

To address the issues of SSTDP, we propose S2-STDP for training a spiking classification layer with one spike per neuron and temporal decision-making. S2-STDP employs error-modulated weight updates with dynamically computed desired timestamps derived from the average firing time within the layer. Then, we propose the PCN training architecture, which further enhances the learning capabilities of our classification layer trained with S2-STDP. PCN encourages neuron specialization through intra-class competition and does not introduce any additional hyperparameters.

#### 4.1 Supervised classification layer

The classification layer is the output layer of our SNN, illustrated in Figure 1C. It is a fully-connected architecture composed of SSIF neurons, designed to process latency-coded spike inputs (i.e. with at most one spike per input, and where the intensity of the input is encoded temporally). For a  $N$ -class problem, the classification layer comprises  $N$  neurons ( $n_1, \dots, n_N$ ), each associated with a distinct class ( $c_1, \dots, c_N$ ). The purpose of the classification layer is to predict the class of a given sample. Since SSIF neurons can fire at most once per sample, we make this prediction based on the output neuron that fires first. This method eliminates the need to propagate the entire input for inference, which can reduce computation time and the number of generated spikes. Formally, the prediction  $\hat{y}$  of the SNN is defined by Equation 6:

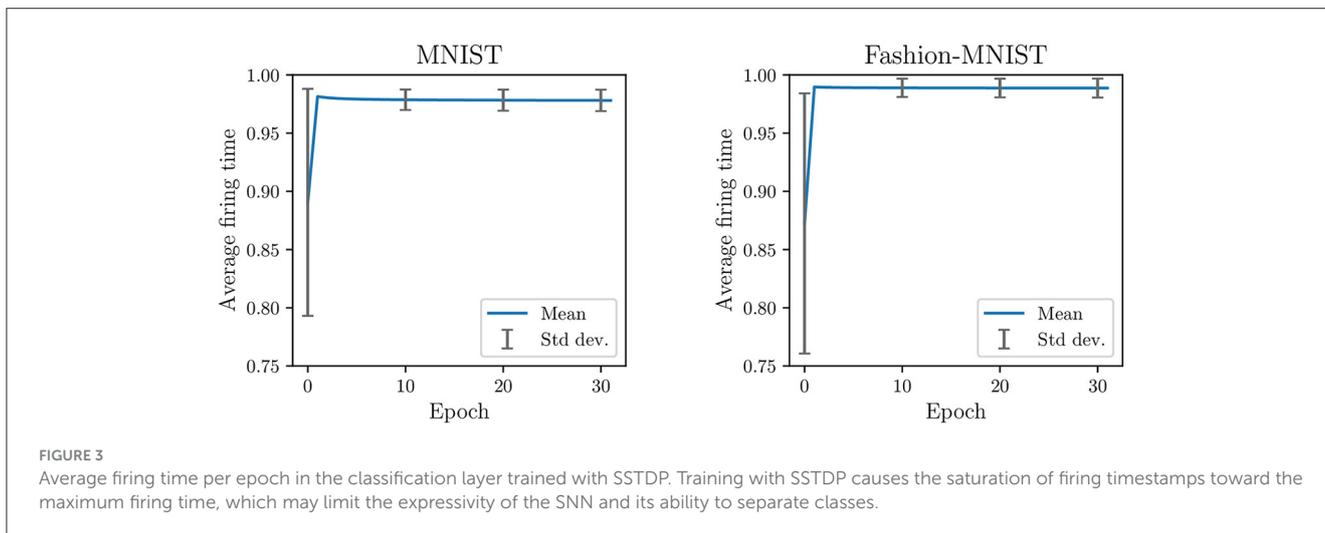
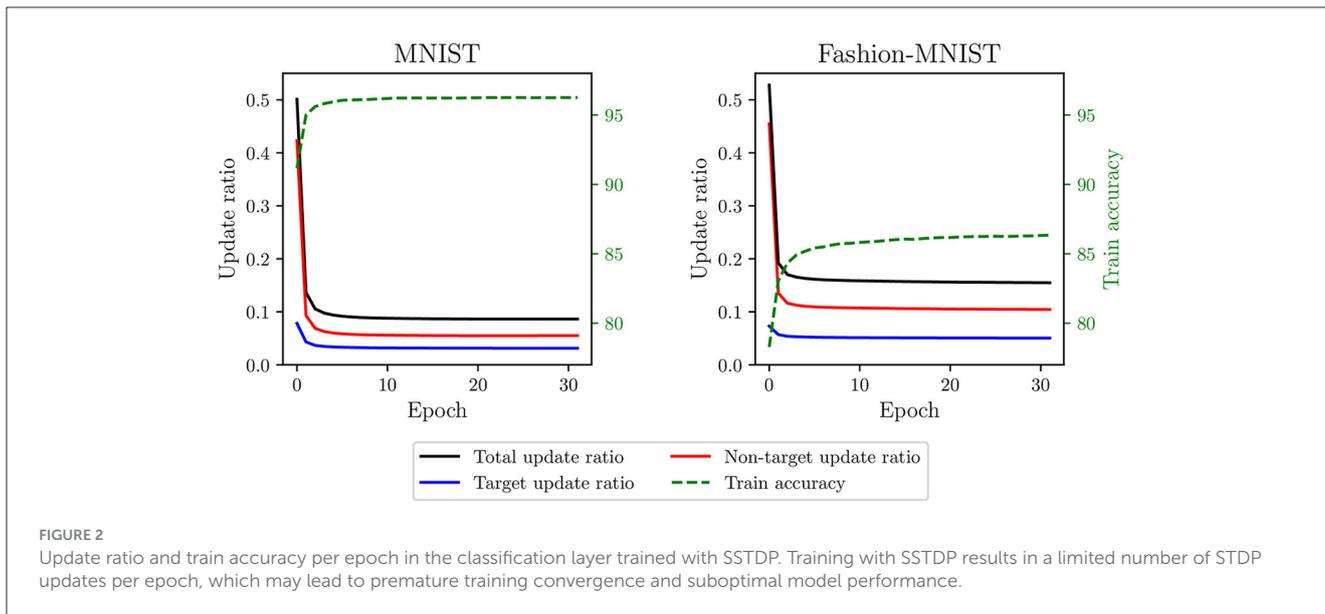
$$\begin{aligned} \hat{y} &= c_{j^*} \\ j^* &= \underset{j \in [1, N]}{\operatorname{argmin}} (t_j), \end{aligned} \quad (6)$$

where  $t_j$  denotes the firing timestamp of neuron  $n_j$ . If several neurons fire at the same timestamp, the one with the highest membrane potential is selected. In practice, the method used for selecting a neuron in the event of a tie has a negligible impact on the performance.

#### 4.2 Stabilized supervised STDP

To optimize the synaptic weights of the classification layer, we propose an error-modulated supervised STDP learning rule named Stabilized Supervised STDP (S2-STDP). This rule teaches neurons

<sup>1</sup> The saturation signifies that the firing timestamps are concentrated toward the upper firing limit, implying an average firing time close to the maximum firing time and a low variance in the firing time.



to alternate their firing between two desired timestamps,  $T_{\text{target}}$  and  $T_{\text{non-target}}$ , dynamically computed for each sample based on  $T_{\text{mean}}$ , the average firing time in the layer. If the class of the neuron corresponds to the class of the sample, the neuron receives a target weight update, teaching it to fire closer to  $T_{\text{target}}$ , just before  $T_{\text{mean}}$ . If the class of the neuron does not correspond to the class of the sample, the neuron receives a non-target weight update, teaching it to fire closer to  $T_{\text{non-target}}$ , right after  $T_{\text{mean}}$ . During training, at the end of the presentation of a sample, the weights of each output neuron are updated with an error-modulated adaptation of the multiplicative STDP following Equation 7:

$$\Delta w_{ij} = \begin{cases} e_j \times A^+ \times \exp\left(-\beta \frac{w_{ij} - w_{\min}}{w_{\max} - w_{\min}}\right) & \text{if } t_j \geq t_i \\ e_j \times A^- \times \exp\left(-\beta \frac{w_{\max} - w_{ij}}{w_{\max} - w_{\min}}\right) & \text{o.w.} \end{cases}, \quad (7)$$

where  $w_{ij}$  is the weight of the synapse connecting input neuron  $n_i$  and output neuron  $n_j$ ,  $\Delta w_{ij}$  is the weight change (such as  $w_{ij} = w_{ij} + \Delta w_{ij}$ ),  $e_j$  is the error of neuron  $n_j$ ,  $A^+$  and  $A^-$

are the positive and negative learning rates,  $\beta$  is the saturation factor,  $w_{\min}$  and  $w_{\max}$  are the minimum and maximum achievable weight values in the layer, and  $t_i$  is the firing timestamp of neuron  $n_i$ . Multiplicative STDP reduces the effect of weight saturation by adjusting the update according to the current weight value and boundaries (Querlioz et al., 2011). Regardless, weights are manually clipped in  $[w_{\min}, w_{\max}]$  after each update to ensure that they remain within a controlled range. The error of an output neuron  $n_j$  is a temporal difference defined by Equation 8:

$$e_j = \frac{t_j - T_j}{T_{\max}}, \quad (8)$$

where  $t_j$  and  $T_j$  respectively represent the actual and desired firing timestamps, and  $T_{\max}$  is the maximum firing time. If an output neuron remains silent during the simulation, we force it to emit a fake spike at  $T_{\max}$ , as in Kheradpisheh and Masquelier (2020). This method ensures that all neurons receive a weight update per sample. To compute the desired firing timestamps, we introduce a

method adapted from SSTDP defined in Equation 9:

$$T_j = \begin{cases} T_{\text{target}} = T_{\text{mean}} - \frac{N-1}{N} \cdot g & \text{if } c_j = y \\ T_{\text{non-target}} = T_{\text{mean}} + \frac{1}{N} \cdot g & \text{if } c_j \neq y \end{cases}, \quad (9)$$

where  $T_{\text{mean}}$  is the average firing time in the layer,  $N$  is the number of neurons,  $c_j$  is the class of neuron  $n_j$ ,  $y$  is the class of the input sample, and  $g$  is a time gap hyperparameter that determines the desired distance from  $T_{\text{mean}}$ . Specifically, whereas SSTDP defines desired time ranges  $[0, T_{\text{target}}]$  and  $[T_{\text{non-target}}, T_{\text{max}}]$ , our adaptation defines desired timestamps  $T_{\text{target}}$  and  $T_{\text{non-target}}$ . Therefore, with our adaptation, neurons can undergo weight updates in both directions, regardless of their associated class. For instance, if a neuron that does match the input class fires before its desired timestamp ( $T_{\text{target}}$ ), its weights will receive a negative update to promote later firing. Conversely, if a neuron that does not match the input class fires after its desired timestamp ( $T_{\text{non-target}}$ ), its weights will receive a positive update to promote earlier firing. With SSTDP, the weights of these neurons would not be updated in such cases.

S2-STDP is carefully designed to resolve the two issues of SSTDP outlined in Section 3. The first aim of S2-STDP is to increase the number of STDP updates per epoch to facilitate training convergence at higher accuracy. This is addressed by defining floating-point desired timestamps rather than time ranges, which are considerably more challenging for neurons to reach. The second aim of S2-STDP is to reduce the saturation of firing timestamps to improve the expressivity of the SNN and its ability to separate classes. This is addressed by enabling positive non-target weight updates that promote earlier firing, and hence, stabilize the output spikes at earlier timestamps. Earlier output firing timestamps may allow the SNN to better fit the specificity of input spikes from a certain class, resulting in a higher variance in the average firing time across samples of different classes. Note that, compared to SSTDP, S2-STDP pushes neurons to fire closer to  $T_{\text{mean}}$ , which may reduce the variance between neuron firing timestamps for a given sample. However, we argue that this is not a concern since the desired timestamps provide more accurate control over the output firing timestamps.

In addition to STDP, we use a heterosynaptic plasticity model (Ferré et al., 2018; Liang et al., 2018) to regulate changes in synaptic weights. This model ensures that all neurons maintain a constant and similar weight average throughout the training process, allowing them equal chances of activation regardless of the number of weight updates they have undergone. After each update of an output neuron  $n_j$ , its weights are normalized following Equation 10:

$$w_{ij} = w_{ij} \cdot \frac{f_{\text{norm}}}{\sum_k w_{kj}}, \quad (10)$$

where  $w_{ij}$  represents the weight of the synapse with input neuron  $n_i$ , and  $f_{\text{norm}}$  is the normalization factor, computed as the sum of weights of neuron  $n_j$  at initialization. In practice, this method also provides robustness against the choice of positive and negative learning rates ( $A^+$  and  $A^-$ ).

## 4.3 Paired competing neurons

S2-STDP involves training each neuron to alternate its firing between two desired timestamps, depending on the class of the input sample (see Equation 9). Hence, the neurons have strong training requirements: they must adapt their weights to satisfy both target and non-target desired timestamps. These requirements can limit their learning capabilities (i.e. their capabilities to learn class-specific patterns) because it is harder to find weight values that allow them to reach both desired timestamps. However, ensuring the convergence toward the desired timestamps is of crucial importance to stabilize the output firing timestamps in the layer. To benefit from the stabilized property of S2-STDP along with enhanced learning capabilities, we propose the Paired Competing Neurons (PCN) training architecture, described in Figure 4.

In this architecture, each class  $c_{j \in [1, N]}$  is associated with a pair of output neurons ( $n_j, n_{j'}$ ) that are connected with lateral inhibition to create intra-class competition. Within each pair and for each sample, the first neuron to fire, called the winner, inhibits the other one, called the loser, and undergoes the weight update. This mechanism is similar to the WTA competition found in STDP-based networks, yet it is class-specific in this context. The purpose of PCN is to encourage, for each class, neuron specialization toward one type of sample: target or non-target. In other words, for a pair of neurons ( $n_j, n_{j'}$ ) associated with class  $c_j$ , when we present samples of class  $y = c_j$ , we want  $n_j$  to fire at  $T_{\text{target}}$  and  $n_{j'}$  to fire after  $n_j$ . Conversely, when we present samples of class  $y \neq c_j$ , we want  $n_{j'}$  to fire at  $T_{\text{non-target}}$  and  $n_j$  to fire after  $n_{j'}$ . The neuron order in this example is arbitrary as we do not assign an objective to each neuron, i.e., within a class-specific pair, we do not explicitly label neurons as target or non-target. Instead, we let this behavior emerge naturally through intra-class competition. Note that the purpose of neurons specializing toward non-target samples extends beyond inhibiting target firing for non-target samples (Tavanaei and Maida, 2015). Here, these neurons also play a crucial role in the S2-STDP training process, as their firing timestamps directly impact the average firing time, and hence, the training convergence. By encouraging specialization toward one type of sample, we reduce training requirements on the neurons because their weights primarily receive one type of update, which improves their learning capabilities.

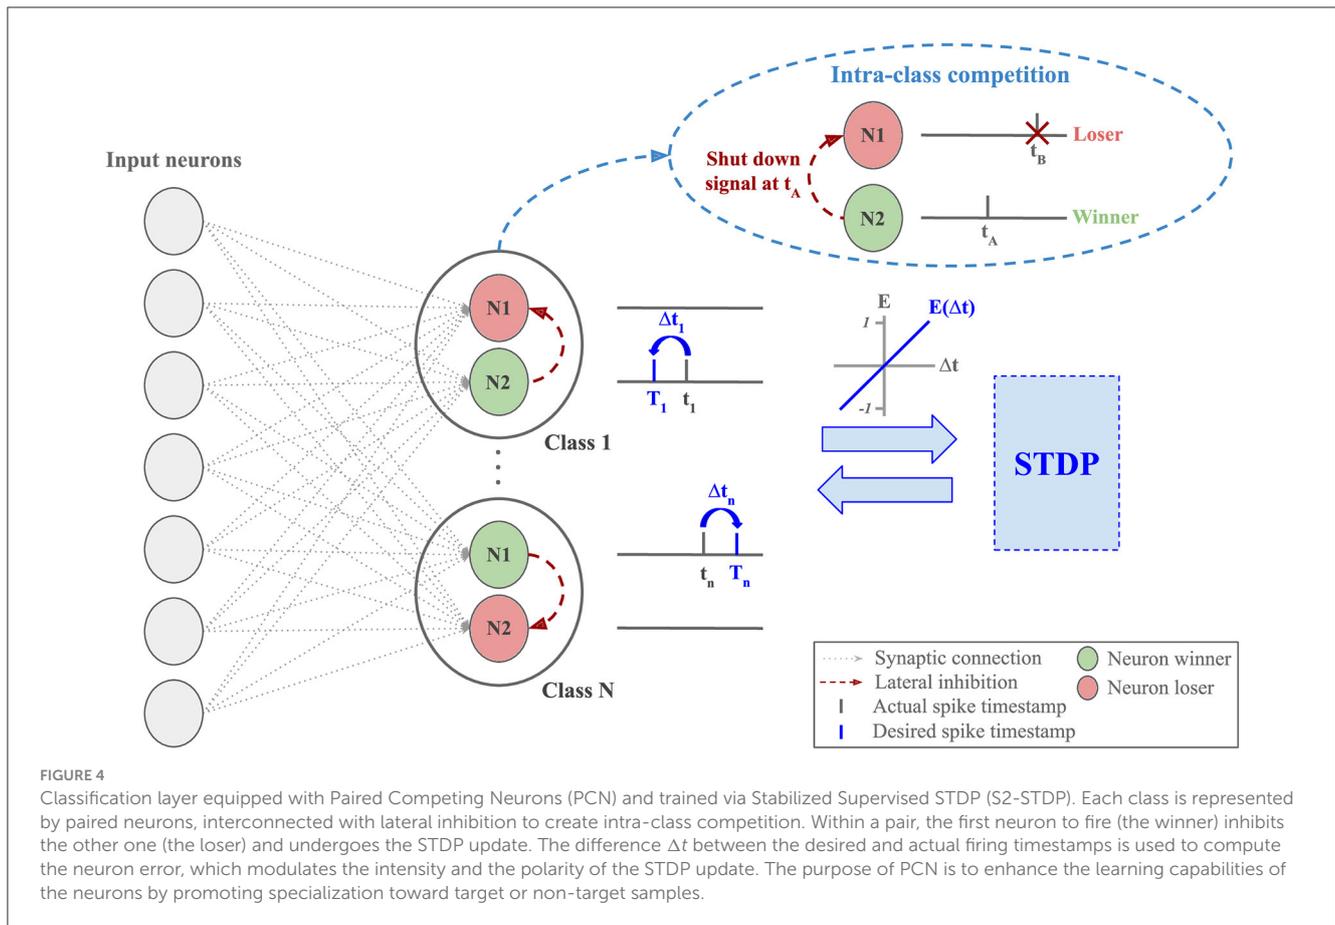
It should be noted that the use of PCN offers notable other advantages. First, thanks to class-wise lateral inhibition, the increased number of neurons does not introduce additional training complexity because, for each sample, only the winners receive weight updates (i.e. one neuron per class). Second, it does not introduce any additional hyperparameters. Third, it can, in principle, be used with any other learning rules involving two desired timestamps.

## 5 Results

### 5.1 Experimental setup

#### 5.1.1 Feature extraction network

Before classification, we extract features from images with a CSNN trained using unsupervised STDP. In our experiments,



we use the CSNN model of Falez et al. (2019b), described in Section 2.3. The employed architecture consists of a single trainable convolutional layer followed by a non-trainable max-pooling layer. Training is done in a layer-wise fashion: the convolutional layer of the CSNN is entirely trained before the training of the classification layer starts. Additional training details are provided in [Supplementary material](#) (Section 1). To analyze the performance of the classification layer across various input sizes, we consider three configurations of the CSNN with increasing numbers of filters: CSNN-16 (16 filters), CSNN-64 (64 filters), and CSNN-128 (128 filters). Within a given dataset, these configurations share the same hyperparameters, except for the number of filters. Unless otherwise specified, the experiments are conducted on CSNN-128.

### 5.1.2 Datasets

We select three image recognition datasets, each comprising ten classes, exhibiting growing complexity: MNIST (LeCun et al., 1998), Fashion-MNIST (Xiao et al., 2017), and CIFAR-10 (Krizhevsky, 2009). Both MNIST and Fashion-MNIST consist of  $28 \times 28$  grayscale images. They contain 60,000 samples for training and 10,000 samples for testing. We preprocess the images with on-center/off-center coding to extract edge information (Falez et al., 2019b). CIFAR-10 is composed of  $32 \times 32$  RGB images, 50,000 for training and 10,000 for testing. We preprocess the images with the hardware-friendly whitening method presented

in Falez et al. (2020) to highlight their edges and high-frequency features. Note that CIFAR-10 is challenging for STDP-based SNNs and only a limited number of studies have considered this dataset thus far (Ferré et al., 2018; Srinivasan and Roy, 2019; Falez et al., 2020; Shrestha et al., 2021). All the preprocessing methods are used with their original hyperparameters, provided in [Supplementary material](#) (Section 2).

### 5.1.3 Protocol

We divide our experimental protocol into two phases: hyperparameter optimization and evaluation. In both phases, we employ an early stopping mechanism (with a patience  $p_{\text{stop}}$ ) during training to prevent overfitting.

During the hyperparameter optimization phase, a subset of the training set is used for validation, which is created by randomly selecting, for each class, a percentage  $p_{\text{val}}$  of its samples. We then apply the gridsearch algorithm to optimize the hyperparameters of the SNN based on the validation accuracy. Hence, we ensure that the hyperparameters are not optimized for the test set. In this work, only the hyperparameters of the classification layer are optimized with gridsearch. The hyperparameters of the CSNN are manually set based on preliminary experiments. For each dataset and model of the classification layer, gridsearch is performed on the CSNN-128 configuration. Then, the same hyperparameters are used with CSNN-64 and CSNN-16, except for the firing threshold that is

divided by 2 and 4, respectively, as the number of input spikes decreases with the input size. All the hyperparameters are provided in [Supplementary material](#) (Section 2).

During the evaluation phase, we use the  $K$ -fold cross-validation strategy. We partition the training set into  $K$  subsets and train  $K$  models, each using a different subset for validation while the remaining  $K - 1$  subsets are used for training. Then, we evaluate the trained models on the test set and we compute the average test accuracy. Note that each model is trained with a different random seed. This allows us to assess the performance of the SNN with varying training and validation data, as well as different weight initializations. In all the following experiments, we choose  $p_{\text{stop}} = 10$ ,  $K = 10$  and  $p_{\text{val}} = \frac{1}{K}$  (i.e. 10% of the training set is used for validation during the hyperparameter optimization phase).

## 5.2 Accuracy comparison

In this section, we present a comparative analysis of the accuracy between two existing methods designed for training SNNs with one spike per neuron, R-STDP and SSTDP (both of which we have implemented and optimized using gridsearch), along with our methods, S2-STDP and S2-STDP+PCN. [Tables 1, 2, 3](#) show the average test accuracy achieved by each method on the MNIST, Fashion-MNIST, and CIFAR-10 datasets, respectively.

Across all datasets and CSNN configurations, S2-STDP consistently outperforms SSTDP. On CSNN-128, we measure a gain of 1.56 pp (percentage points) for MNIST, 0.72 pp for Fashion-MNIST, and 0.65 pp for CIFAR-10. In addition, SSTDP tends to underperform when confronted with smaller input sizes. For instance, on CSNN-16, S2-STDP outperforms SSTDP by 10.91 pp on MNIST and 10.16 pp on CIFAR-10. The lower performance of SSTDP on the CSNN-16 configuration of these datasets is caused by training divergence, as evidenced by the higher standard deviations, and is influenced by the employed hyperparameters (transferred from CSNN-128). It suggests that SSTDP lacks robustness against hyperparameters. While the accuracy gain between SSTDP and S2-STDP is not always substantial, our adaptation enables more effective training of the classification layer, irrespective of the number of input features. More importantly, S2-STDP leverages compatibility with the PCN architecture, which further improves the accuracy of S2-STDP across all datasets and CSNN configurations, without requiring any additional hyperparameters. When integrating PCN with S2-STDP on CSNN-128, we measure an additional gain of 0.78 pp for MNIST, 1.24 pp for Fashion-MNIST, and 1.28 pp for CIFAR-10. In comparison to the other existing STDP-based methods, S2-STDP+PCN achieves the highest accuracy on the MNIST and Fashion-MNIST datasets. Specifically, compared to SSTDP on CSNN-128, it shows an accuracy improvement of 2.34 pp on MNIST, 1.96 pp on Fashion-MNIST, and 1.93 pp on CIFAR-10. We recall that our proposed training methods based on S2-STDP employ a weight normalization mechanism not employed by SSTDP. We evaluate the effect of weight normalization in [Supplementary material](#) (Section 3.3).

On CIFAR-10, R-STDP significantly outperforms all the STDP-based methods. However, R-STDP requires 200 output neurons

to achieve this performance whereas S2-STDP+PCN only uses 20 neurons. When R-STDP is used with 20 output neurons as S2-STDP+PCN, it performs significantly worse than all other methods across all datasets and CSNN configurations. This observation highlights the importance of error-modulated weight updates in enabling effective and efficient supervised training with STDP. In our case, it leads to a substantial reduction in the number of trainable parameters by a factor of 10. Also, on Fashion-MNIST, R-STDP obtains relatively low performance and fails to extract more relevant features when the number of feature maps increases. The accuracy gain between CSNN-16 and CSNN-128 with R-STDP is only about 0.96 pp, whereas it is about 3.23 pp with S2-STDP+PCN.

## 5.3 S2-STDP addresses the issues of SSTDP

In Section 3, we elaborated on the issues of SSTDP: the limited number of STDP updates per epoch and the saturation of firing timestamps toward the maximum firing time. Our proposed S2-STDP is specifically designed to address these issues by defining desired timestamps that stabilize the output spikes at earlier timestamps. In this section, we analyze the effect of S2-STDP on these issues.

In [Figure 5](#), we compare the update ratio, average firing time, and train accuracy per epoch in the classification layer trained with the various SSTDP-based methods on the MNIST dataset. As previously illustrated in a preliminary experiment (see [Figures 2, 3](#)), SSTDP suffers from a limited number of updates per epoch, firing timestamps close to the maximum firing time, and low variance in the average firing time between samples. Our proposed methods based on S2-STDP successfully address these three issues. In comparison to SSTDP at epoch 30, S2-STDP increases the update ratio from 9% to nearly 100%, reduces the average firing time from 0.98 to 0.88, and augments its standard deviation from 0.009 to 0.05. Note that with S2-STDP+PCN, the update ratio is close to 50% (instead of 100%), which is the maximum achievable value because half of the neurons are inhibited. By addressing the issues of SSTDP, our methods based on S2-STDP enable training convergence at higher accuracies. In addition, due to the increased number of updates per epoch, they achieve higher accuracies in fewer epochs. For instance, S2-STDP reaches a training accuracy of 95.56% after the first epoch, whereas SSTDP only reaches 91.16% and requires two additional epochs to reach 95.60%. In [Supplementary material](#) (Section 3.1), we show similar results for the Fashion-MNIST dataset and we demonstrate that the resolution of these issues is not attributed to the additional weight normalization mechanism employed.

We mentioned in Section 4.2 that S2-STDP pushes neurons to fire closer to the average firing time, and hence, to each other, compared to SSTDP. In the following experiment, we show that contrary to intuition, a tighter distribution of output firing timestamps does not necessarily reduce the ability of the SNN to separate classes. [Figure 6](#) illustrates the distribution of firing time differences between the first non-target neuron to fire and the target neuron, on MNIST test samples. Since the SNN prediction is based on the neuron that fires first, the sign of the time difference indicates the ability of the SNN to classify the sample: a positive time

TABLE 1 Accuracy of the existing and proposed learning rules on MNIST.

Rule	Output neurons	CSNN-16	CSNN-64	CSNN-128
R-STDP	200	96.05 ± 0.51	97.47 ± 0.16	97.88 ± 0.13
R-STDP	20	61.23 ± 27.74	93.34 ± 0.66	93.28 ± 0.87
SSTDP	10	84.49 ± 7.07	95.56 ± 0.13	96.25 ± 0.11
S2-STDP*	10	95.40 ± 0.21	97.19 ± 0.09	97.81 ± 0.05
S2-STDP+PCN*	20	97.02 ± 0.14	98.21 ± 0.07	98.59 ± 0.06

\* Ours.

TABLE 2 Accuracy of the existing and proposed learning rules on Fashion-MNIST.

Rule	Output neurons	CSNN-16	CSNN-64	CSNN-128
R-STDP	200	82.30 ± 0.26	82.30 ± 0.92	83.26 ± 0.22
R-STDP	20	76.50 ± 0.52	76.73 ± 0.13	77.01 ± 0.22
SSTDP	10	81.40 ± 0.67	84.51 ± 0.18	85.16 ± 0.10
S2-STDP*	10	82.23 ± 0.40	84.92 ± 0.24	85.88 ± 0.22
S2-STDP+PCN*	20	83.89 ± 0.40	85.84 ± 0.19	87.12 ± 0.21

\* Ours.

TABLE 3 Accuracy of the existing and proposed learning rules on CIFAR-10.

Rule	Output Neurons	CSNN-16	CSNN-64	CSNN-128
R-STDP	200	51.55 ± 1.23	61.85 ± 0.61	65.56 ± 0.38
R-STDP	20	38.51 ± 4.41	51.74 ± 0.93	54.02 ± 0.80
SSTDP	10	37.78 ± 2.41	57.53 ± 0.38	60.88 ± 0.23
S2-STDP*	10	47.94 ± 0.49	58.24 ± 0.27	61.53 ± 0.16
S2-STDP+PCN*	20	49.23 ± 0.56	59.58 ± 0.16	62.81 ± 0.15

\* Ours.

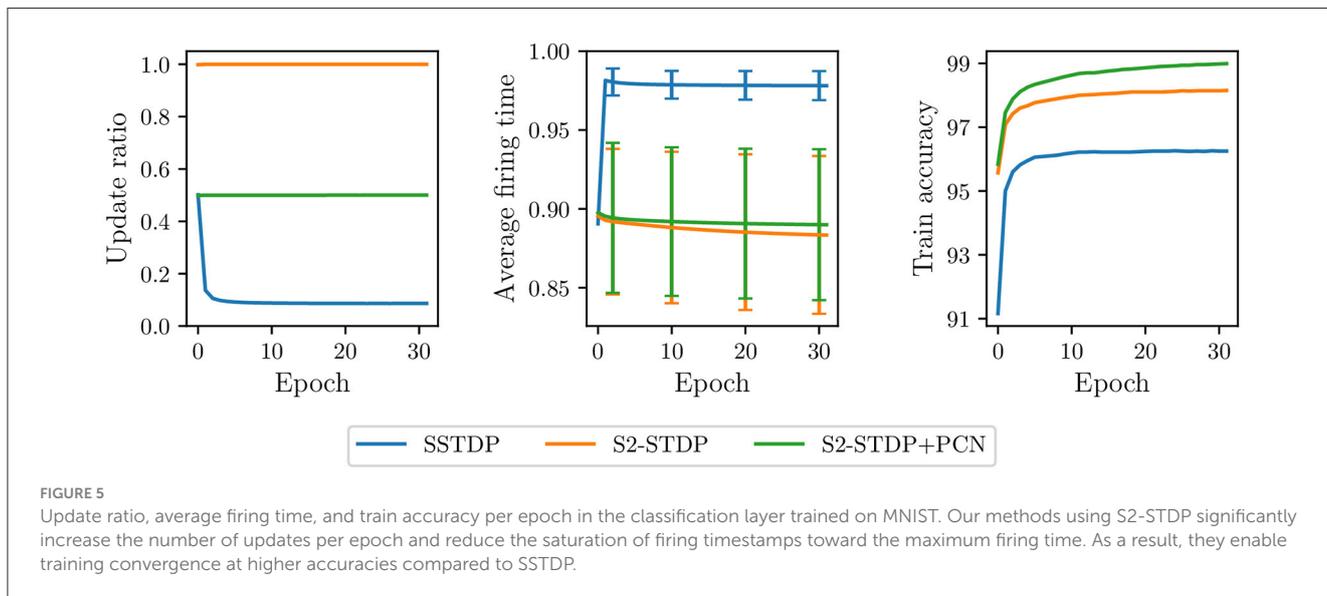


FIGURE 5

Update ratio, average firing time, and train accuracy per epoch in the classification layer trained on MNIST. Our methods using S2-STDP significantly increase the number of updates per epoch and reduce the saturation of firing timestamps toward the maximum firing time. As a result, they enable training convergence at higher accuracies compared to SSTDTP.

difference indicates a correctly classified sample (the target neuron fires before the non-target neuron), and conversely. The tight distribution of S2-STDP implies that the firing time differences tend to be significantly smaller compared to SSTDTP. This behavior arises

from our adapted training process, which pushes neurons to fire closer to the average firing time. While it may seem intuitive to maximize the firing time difference for improved class separability, a closer examination of negative time differences reveals that

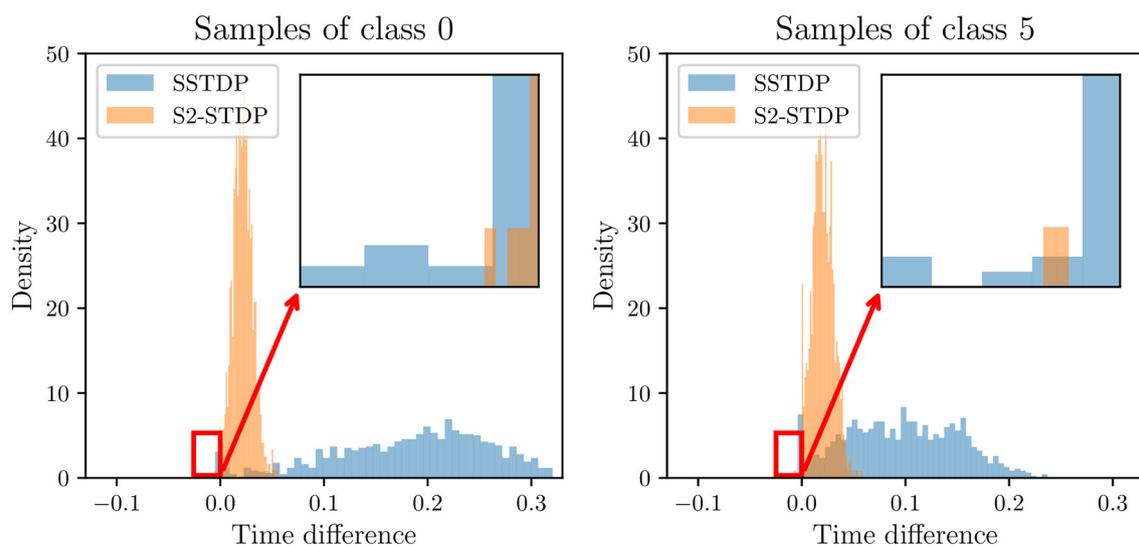


FIGURE 6

Distribution of firing time differences between the first non-target neuron to fire and the target neuron, in the classification layer for MNIST test samples of class 0 and 5. A negative time difference indicates that the target neuron fired after a non-target neuron, leading to misclassification of the sample. S2-STDP achieves a tighter distribution of firing time differences but results in fewer misclassified samples compared to SSTDP.

SSTDP leads to more misclassified samples than S2-STDP (which is confirmed by the overall lower accuracy of SSTDP). In the context of temporal decision-making, we argue that supervised STDP rules should not necessarily aim to maximize the firing time difference between the target and non-target neurons. Instead, it seems more important to ensure accurate control over the output firing timestamps. In [Supplementary material](#) (Section 3.1), we present similar results with the other classes of the MNIST dataset.

#### 5.4 PCN enables neuron specialization

The PCN architecture exploits the two desired timestamps defined by S2-STDP along with intra-class competition to promote neuron specialization toward one type of sample: target or non-target. In this section, we study the impact of integrating a PCN architecture on the output firing timestamps of our classification layer trained with S2-STDP.

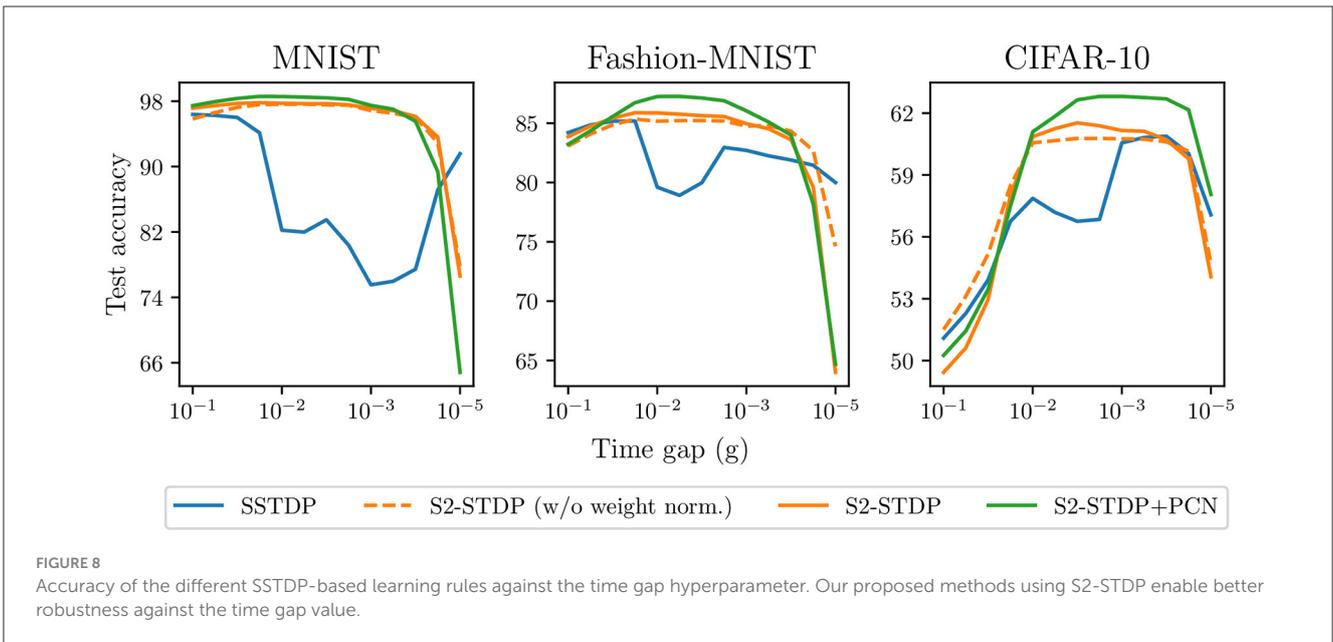
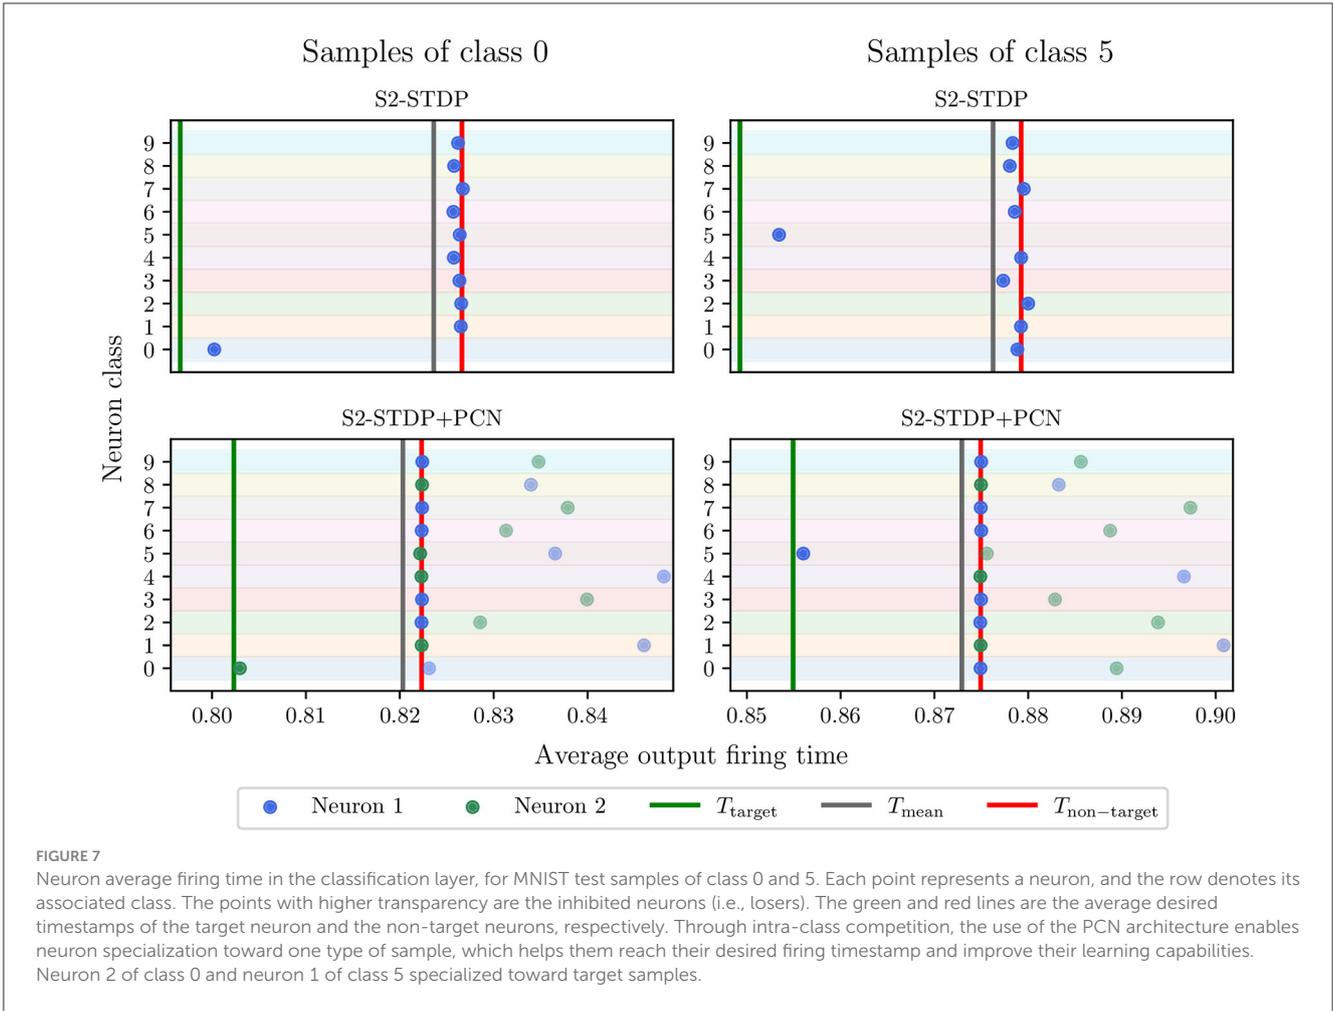
[Figure 7](#) illustrates the average firing time of output neurons trained with and without PCN, on MNIST test samples of class 0 and 5. Note that the time gaps used for S2-STDP and S2-STDP+PCN differ but they correspond to the optimal value obtained through gridsearch. We observe that neurons trained with a PCN architecture better reach their desired timestamps, particularly for the non-target neurons (see Section 3.2 of [Supplementary material](#) for additional analysis). This is because, through intra-class competition, PCN naturally enables neuron specialization toward target or non-target samples. This specialization is illustrated by neuron 2 of class 0 firing at  $T_{\text{target}}$  when presented to samples of class 0 (target class) and being inhibited by neuron 1 of class 0 for samples of class 5 (non-target class). Conversely, neuron 2 of class 5 fires at  $T_{\text{non-target}}$  when presented to samples of class 0 (non-target class) and is inhibited

by neuron 1 of class 5 for samples of class 5 (target class). Hence, the emerging target neuron may differ between classes (neuron 1 or 2). Intra-class competition reduces training requirements on neurons specialized toward target samples because they do not have to fire at an exact desired timestamp ( $T_{\text{non-target}}$ ) for non-target samples. As observed, inhibited neurons for non-target samples (i.e. neurons specialized toward target samples) have much larger firing timestamps compared to the non-target desired timestamp. Consequently, by creating neuron specialization toward target or non-target samples, PCN facilitates the learning of more specific patterns, which improves their learning capabilities.

#### 5.5 Robustness against the time gap hyperparameter

During training, the time gap hyperparameter  $g$  of S2-STDP is used to define the distance between the desired timestamps and the average firing time. Selecting an appropriate value for this hyperparameter is crucial to ensure accurate class separation. However, hyperparameter tuning can be a time-consuming task, and achieving the optimal value may not always be feasible. In this section, we investigate the influence of the time gap value on accuracy.

[Figure 8](#) compares SSTDP and our proposed methods across various values of  $g$  on the three datasets. S2-STDP demonstrates greater robustness compared to SSTDP regarding the choice of  $g$ , as it significantly expands the range of values that can achieve near-optimal accuracy. The accuracy curve of S2-STDP exhibits a more pronounced bell-shaped pattern with a larger plateau near the maximum. This implies that tuning  $g$  can be easier with S2-STDP. When considering a suitable range for  $g$ , S2-STDP always achieves higher or



similar accuracy compared to SSTDP. We also illustrate the accuracy of S2-STDP without weight normalization to demonstrate that its improved robustness is not due to

this additional mechanism. The use of PCN as a training architecture for S2-STDP almost always improves its performance. Hence, the efficacy of PCN is not dependent on a specific

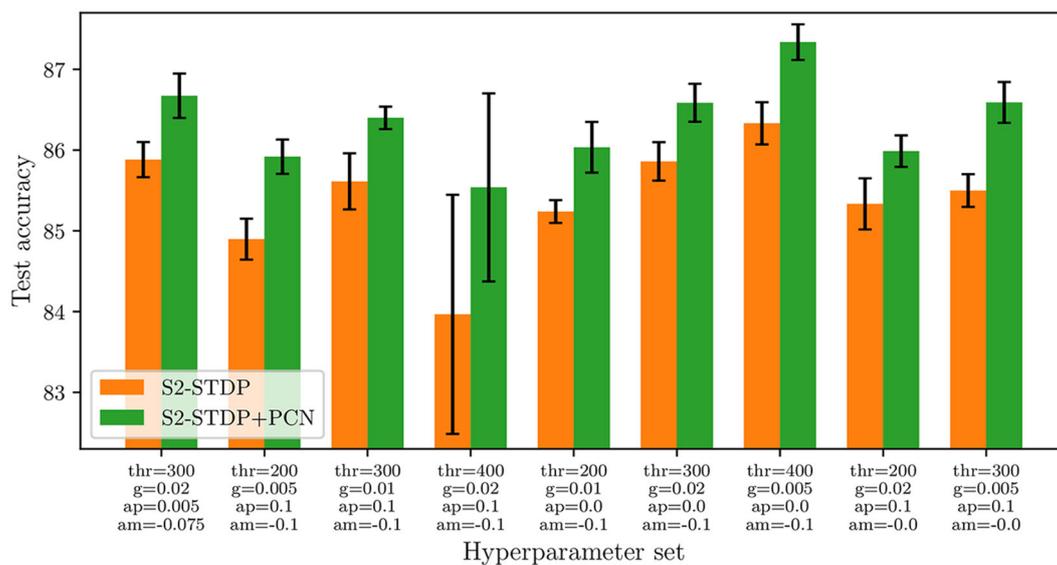


FIGURE 9

Accuracy of S2-STDP, with and without the use of the PCN architecture, across different hyperparameter sets on Fashion-MNIST. PCN always improves S2-STDP performance, without introducing any additional hyperparameters.

value of  $g$ . Note that all the methods use their respective gridsearch-optimized hyperparameters.

## 5.6 Robustness against the hyperparameter set

Our PCN training architecture improves the accuracy of S2-STDP across all evaluated datasets, as indicated in Section 5.2. However, the hyperparameters of the classification layer are individually optimized for S2-STDP and S2-STDP+PCN, suggesting that PCN might only enhance performance with specific hyperparameter sets. In this section, we present evidence that PCN can effectively improve the accuracy of S2-STDP irrespective of the hyperparameters used. Specifically, we aim to demonstrate that integrating PCN into an existing SNN employing S2-STDP for training is likely to always yield improved results.

Figure 9 illustrates the comparison of S2-STDP accuracy with and without PCN across different hyperparameter sets on the Fashion-MNIST dataset. Results on the other datasets are reported in Supplementary material (Section 3.4). The varying hyperparameters include the firing threshold ( $thr$ ), time gap ( $g$ ), and learning rates ( $ap$  and  $am$ ). Their values are selected within a suitable range to achieve satisfactory performance with S2-STDP. The results show that integrating PCN improves the accuracy of S2-STDP consistently, with an average improvement of 0.94 pp and a maximum improvement of 1.57 pp. For comparison, when independently optimizing the hyperparameters for both methods, the measured accuracy improvement is 1.24 pp. Hence, the PCN architecture is an effective method for easily improving the performance of a spiking classifier trained with S2-STDP, regardless of the hyperparameter set and without introducing any additional hyperparameters. As detailed in Supplementary material, the

results on the other datasets exhibit consistency with the analysis conducted on Fashion-MNIST.

## 5.7 Comparison with the literature

In Table 4, we present an accuracy comparison between our partially supervised SNN, trained with STDP and S2-STDP+PCN, along with other existing algorithms employed for training SNNs. For this comparison, we focused on supervised methods, primarily STDP-based and BP-based with local updates. Note that the reported accuracies come from the original papers.

On the MNIST and Fashion-MNIST datasets, our proposed SNN outperforms the partially supervised approaches (Mozafari et al., 2019; Hao et al., 2020) and demonstrates competitive performance with most of the fully supervised approaches. It is important to mention that, unlike these methods, only the output layer of our SNN is trained with supervision. For instance, on Fashion-MNIST, Zhao et al. (2020) employ a 6-layer SNN trained with global feedback + STDP, achieving an accuracy of 89.05%. In contrast, we employ a 2-layer partially supervised CSNN, resulting in an accuracy loss of only 1.93 pp. Additionally, our results are highly competitive with a 2-layer SNN trained using a BP-based algorithm, with an accuracy loss of 0.88 pp.

On the CIFAR-10 dataset, the performance of our method remains significantly low compared to state-of-the-art algorithms. However, these algorithms, like the VGG-7 trained with SSTDP (Liu et al., 2021b) (which reports a top-1 accuracy and not an average), rely on non-local learning rules that cannot be employed for direct training on neuromorphic hardware. On the contrary, our proposed SNN is exclusively trained using local learning rules and restricts the use of supervision to one layer. In comparison to other local-based methods, our SNN demonstrates

TABLE 4 Accuracy comparison of our proposed SNN with the literature.

Dataset	References	Number of layers unsupervised / supervised	Learning rule	Local	Accuracy (%)
MNIST	Lee et al., 2018	0 / 4	<b>STDP pretraining + BP</b>	<b>no</b>	<b>99.28</b>
	Liu et al., 2021b	0 / 2	SSTDP	no	98.10
	Tavanaei and Maida, 2019	0 / 3	BP-STDP	no	97.20
	Zhang et al., 2021	0 / 3	BRP	yes	99.01
	Shrestha et al., 2021	0 / 4	EMSTDP	yes	98.90
	Mozafari et al., 2019	2 / 1	STDP + R-STDP	yes	97.20
	Hao et al., 2020	1 / 1	Sym-STDP	yes	96.73
	This work	1 / 1	STDP + S2-STDP+PCN	yes	98.59
Fashion-MNIST	Kheradpisheh and Masquelier, 2020	0 / 2	BP	no	88.00
	Mirsadeghi et al., 2023	0 / 4	<b>STiDi-BP</b>	<b>yes</b>	<b>92.80</b>
	Shrestha et al., 2021	0 / 4	EMSTDP	yes	92.70
	Zhao et al., 2020	0 / 6	Global feedback + STDP	yes	89.05
	Hao et al., 2020	1 / 1	Sym-STDP	yes	85.31
	This work	1 / 1	STDP + S2-STDP+PCN	yes	87.12
CIFAR-10	Liu et al., 2021b	0 / 7	<b>SSTDP</b>	<b>no</b>	<b>91.31</b>
	Ferré et al., 2018	1 / 3	STDP + ANN-based BP	no	71.20
	Srinivasan and Roy, 2019	1 / 1	STDP + ANN-based BP	no	66.23
	Shrestha et al., 2021	0 / 4	EMSTDP	yes	64.40
	Zhang et al., 2021	0 / 3	BRP	yes	57.08
	This work	1 / 1	STDP + S2-STDP+PCN	yes	62.81

Accuracies come from the original papers. Bold corresponds to the work with the highest accuracy.

competitive results. Overall, this dataset is particularly challenging for shallow architectures with only one supervised layer. We believe that the features extracted by our unsupervised layer are not distinguishable enough to allow for an accurate analysis.

## 6 Conclusion

In this paper, we proposed Stabilized Supervised STDP (S2-STDP), a supervised STDP learning rule for training a spiking classification layer with one spike per neuron and temporal decision-making. This layer can be employed to classify features extracted by a convolutional SNN (CSNN) equipped with unsupervised STDP. Our learning rule integrates error-modulated weight updates that align neuron spikes with desired timestamps derived from the average firing time within the layer. Then, to further enhance the learning capabilities of the classification layer trained with S2-STDP, we introduced a training architecture called Paired Competing Neurons (PCN). PCN associates each class with paired neurons connected via lateral inhibition and encourages neuron specialization through intra-class competition. We evaluated S2-STDP and PCN on three image recognition datasets of growing complexity: MNIST, Fashion-MNIST, and CIFAR-10. Experiments showed that our methods outperform state-of-the-art supervised STDP rules when employed to train our spiking classification layer. S2-STDP successfully addresses

the issues of SSTDP concerning the limited number of STDP updates per epoch and the saturation of firing timestamps toward the maximum firing time. The PCN architecture enhances the performance of S2-STDP, regardless of the hyperparameter set and without introducing any additional hyperparameters. Our methods also exhibited improved hyperparameter robustness as compared to SSTDP.

In the future, we plan to expand S2-STDP to multi-layer architectures, while maintaining the local computation required for on-chip learning. This includes exploring both feedback connections (Zhao et al., 2020) and local losses (Mirsadeghi et al., 2021).

## Data availability statement

The original contributions presented in the study are included in the article/Supplementary material, further inquiries can be directed to the corresponding author.

## Author contributions

GG: Writing – original draft, Writing – review & editing, Conceptualization, Methodology, Software, Validation, Visualization. PT: Writing – review & editing, Methodology,

Supervision. IB: Writing – review & editing, Funding acquisition, Methodology, Supervision.

## Funding

The author(s) declare that financial support was received for the research, authorship, and/or publication of this article. This work is funded by Chaire Luxant-ANVI (Métropole Européenne de Lille) and supported by IRCICA (CNRS UAR 3380).

## Acknowledgments

We would like to thank Benoit Miramond for the helpful exchange.

## Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships

## References

- Abiodun, O. I., Jantan, A., Omolara, A. E., Dada, K. V., Mohamed, N. A., and Arshad, H. (2018). State-of-the-art in artificial neural network applications: a survey. *Heliyon* 4:e00938. doi: 10.1016/j.heliyon.2018.e00938
- Bengio, Y., Courville, A., and Vincent, P. (2013). Representation learning: a review and new perspectives. *IEEE Trans. Pattern Anal. Mach. Intell.* 35, 1798–1828. doi: 10.1109/TPAMI.2013.50
- Caporale, N., and Dan, Y. (2008). Spike timing-dependent plasticity: a Hebbian learning rule. *Annu. Rev. Neurosci.* 31, 25–46. doi: 10.1146/annurev.neuro.31.060407.125639
- Dampfhofer, M., Mesquida, T., Valentian, A., and Anghel, L. (2023). Backpropagation-based learning techniques for deep spiking neural networks: a survey. *IEEE Trans. Neural. Netw. Learn. Syst.* doi: 10.1109/TNNLS.2023.3263008
- Eshraghian, J. K., Ward, M., Neftci, E., Wang, X., Lenz, G., Dwivedi, G., et al. (2021). Training spiking neural networks using lessons from deep learning. *ArXiv [preprint]arXiv:2109.12894*.
- Falez, P., Tirilly, P., and Bilasco, I. M. (2020). “Improving STDP-based visual feature learning with whitening,” in *International Joint Conference on Neural Networks (IEEE)*.
- Falez, P., Tirilly, P., Bilasco, I. M., Devienne, P., and Boulet, P. (2019a). Unsupervised visual feature learning with sparsely activated convolutional spiking neural networks enabled by voltage-dependent synaptic plasticity. *Neurom. Comp. Eng.* 3:acad98. doi: 10.1088/2634-4386/acad98
- Falez, P., Tirilly, P., Marius Bilasco, I., Devienne, P., and Boulet, P. (2019b). “Multi-layered spiking neural network with target timestamp threshold adaptation and STDP,” in *International Joint Conference on Neural Networks (IEEE)*.
- Ferré, P., Mamalet, F., and Thorpe, S. J. (2018). Unsupervised feature learning with winner-takes-all based STDP. *Front. Comput. Neurosci.* 12:24. doi: 10.3389/fncom.2018.00024
- Frémaux, N., and Gerstner, W. (2015). Neuromodulated spike-timing-dependent plasticity, and theory of three-factor learning rules. *Front. Neural Circuits* 9:85. doi: 10.3389/fncir.2015.00085
- Goupy, G., Juneau-Fecteau, A., Garg, N., Balafrej, I., Alibart, F., Frechette, L., et al. (2023). Unsupervised and efficient learning in sparsely activated convolutional spiking neural networks enabled by voltage-dependent synaptic plasticity. *Neurom. Comp. Eng.* 3:acad98. doi: 10.1088/2634-4386/acad98
- Guo, W., Fouda, M. E., Eltawil, A. M., and Salama, K. N. (2021). Neural coding in spiking neural networks: a comparative study for robust neuromorphic systems. *Front. Neurosci.* 15:638474. doi: 10.3389/fnins.2021.638474
- Hao, Y., Huang, X., Dong, M., and Xu, B. (2020). A biologically plausible supervised learning method for spiking neural networks using the symmetric STDP rule. *Neural Networks* 121, 387–395. doi: 10.1016/j.neunet.2019.09.007
- Hebb, D. (1949). *The Organization of Behavior*. Berlin, Heidelberg: Springer.
- Jeong, D. S., Kim, K. M., Kim, S., Choi, B. J., and Hwang, C. S. (2016). Memristors for energy-efficient new computing paradigms. *Adv. Elect. Mat.* 2:90. doi: 10.1002/aelm.201600090
- Khacef, L., Klein, P., Cartiglia, M., Rubino, A., Indiveri, G., and Chicca, E. (2023). Spike-based local synaptic plasticity: a survey of computational models and neuromorphic circuits. *Neurom. Comp. Eng.* 3:ad05da. doi: 10.1088/2634-4386/ad05da
- 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
- Kheradpisheh, S. R., and Masquelier, T. (2020). Temporal backpropagation for spiking neural networks with one spike per neuron. *Int. J. Neural Syst.* 30:276. doi: 10.1142/S0129065720500276
- Krizhevsky, A. (2009). *Learning Multiple Layers of Features from Tiny Images*. Toronto, ON: University of Toronto.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 2278–2323. doi: 10.1109/5.726791
- Lee, C., Panda, P., Srinivasan, G., and Roy, K. (2018). Training deep spiking convolutional neural networks with STDP-based unsupervised pre-training followed by supervised fine-tuning. *Front. Neurosci.* 12:435. doi: 10.3389/fnins.2018.00435
- Lee, C., Srinivasan, G., Panda, P., and Roy, K. (2019). Deep spiking convolutional neural network trained with unsupervised spike-timing-dependent plasticity. *Trans. Cognit. Dev. Syst.* 11, 384–394. doi: 10.1109/TCDS.2018.2833071
- Li, J., Xu, H., Sun, S.-Y., Li, N., Li, Q., Li, Z., et al. (2022). *In situ* learning in hardware compatible multilayer memristive spiking neural network. *Trans. Cognit. Dev. Syst.* 14, 448–461. doi: 10.1109/TCDS.2021.3049487
- Liang, Z., Schwartz, D., Ditzler, G., and Koyluoglu, O. O. (2018). The impact of encoding-decoding schemes and weight normalization in spiking neural networks. *Neural Netw.* 108, 365–378. doi: 10.1016/j.neunet.2018.08.024
- Liu, D., Yu, H., and Chai, Y. (2021a). Low-power computing with neuromorphic engineering. *Adv. Intellig. Syst.* 3:150. doi: 10.1002/aisy.202000150
- Liu, F., Zhao, W., Chen, Y., Wang, Z., Yang, T., and Jiang, L. (2021b). SSTDP: supervised spike timing dependent plasticity for efficient spiking neural network training. *Front. Neurosci.* 15:756876. doi: 10.3389/fnins.2021.756876
- Ma, C., Xu, J., and Yu, Q. (2021). “Temporal dependent local learning for deep spiking neural networks,” in *International Joint Conference on Neural Networks (IEEE)*.
- Milo, V., Malavena, G., Monzio Compagnoni, C., and Ielmini, D. (2020). Memristive and CMOS devices for neuromorphic computing. *Materials* 13:166. doi: 10.3390/ma13010166

that could be construed as a potential conflict of interest.

## Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

## Supplementary material

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fnins.2024.1401690/full#supplementary-material>

- Mirsadeghi, M., Shalchian, M., Kheradpisheh, S. R., and Masquelier, T. (2021). STiDi-BP: Spike time displacement based error backpropagation in multilayer spiking neural networks. *Neurocomputing* 427, 131–140. doi: 10.1016/j.neucom.2020.11.052
- Mirsadeghi, M., Shalchian, M., Kheradpisheh, S. R., and Masquelier, T. (2023). Spike time displacement-based error backpropagation in convolutional spiking neural networks. *Neural Comp. Appl.* 35, 15891–15906. doi: 10.1007/s00521-023-08567-0
- Mozafari, M., Ganjtabesh, M., Nowzari-Dalini, A., Thorpe, S. J., and Masquelier, T. (2019). Bio-inspired digit recognition using reward-modulated spike-timing-dependent plasticity in deep convolutional networks. *Pattern Recognit.* 94:15. doi: 10.1016/j.patcog.2019.05.015
- Neftci, E. O., Augustine, C., Paul, S., and Deterakis, G. (2017). Event-driven random back-propagation: enabling neuromorphic deep learning machines. *Front. Neurosci.* 11:324. doi: 10.3389/fnins.2017.00324
- Park, S., Kim, S., Na, B., and Yoon, S. (2020). “T2FSNN: deep spiking neural networks with time-to-first-spike coding,” in *Design Automation Conference*.
- Ponulak, F., and Kasinski, A. (2010). Supervised learning in spiking neural networks with ReSuMe: sequence learning, classification, and spike shifting. *Neural Comput.* 22, 467–510. doi: 10.1162/neco.2009.11-08-901
- Ponulak, F., and Kasinski, A. (2011). Introduction to spiking neural networks: information processing, learning and applications. *Acta Neurobiol. Exp.* 71, 409–433. doi: 10.55782/ane-2011-1862
- Querlioz, D., Bichler, O., and Gamrat, C. (2011). “Simulation of a memristor-based spiking neural network immune to device variations,” in *International Joint Conference on Neural Networks*, 1775–1781 (IEEE).
- Rullen, R. V., and Thorpe, S. J. (2001). Rate coding versus temporal order coding: what the retinal ganglion cells tell the visual cortex. *Neural Comput.* 13, 1255–1283. doi: 10.1162/08997660152002852
- Saighi, S., Mayr, C. G., Serrano-Gotarredona, T., Schmidt, H., Lecerf, G., Tomas, J., et al. (2015). Plasticity in memristive devices for spiking neural networks. *Front. Neurosci.* 9:51. doi: 10.3389/fnins.2015.00051
- Saranirad, V., Dora, S., McGinnity, T. M., and Coyle, D. (2022). “Assembly-based STDP: a new learning rule for spiking neural networks inspired by biological assemblies,” in *International Joint Conference on Neural Networks* (IEEE).
- Schuman, C. D., Potok, T. E., Patton, R. M., Birdwell, J. D., Dean, M. E., Rose, G. S., et al. (2017). A survey of neuromorphic computing and neural networks in hardware. *arXiv*. [preprint]arXiv:1705.06963. doi: 10.48550/arXiv.1705.06963
- Shrestha, A., Ahmed, K., Wang, Y., and Qiu, Q. (2017). “Stable spike-timing dependent plasticity rule for multilayer unsupervised and supervised learning,” in *International Joint Conference on Neural Networks*, 1999–2006.
- Shrestha, A., Fang, H., Mei, Z., Rider, D. P., Wu, Q., and Qiu, Q. (2022). A survey on neuromorphic computing: models and hardware. *Circuits Syst. s Magaz.* 22, 6–35. doi: 10.1109/MCAS.2022.3166331
- Shrestha, A., Fang, H., Rider, D. P., Mei, Z., and Qiu, Q. (2021). “In-hardware learning of multilayer spiking neural networks on a neuromorphic processor,” in *Design Automation Conference* (ACM/IEEE), 367–372.
- Shrestha, A., Fang, H., Wu, Q., and Qiu, Q. (2019). “Approximating back-propagation for a biologically plausible local learning rule in spiking neural networks,” in *International Conference on Neuromorphic Systems* (ACM).
- Srinivasan, G., and Roy, K. (2019). ReStoCNet: residual stochastic binary convolutional spiking neural network for memory-efficient neuromorphic computing. *Front. Neurosci.* 13:189. doi: 10.3389/fnins.2019.00189
- Tavanaei, A., and Maida, A. (2019). BP-STDP: approximating backpropagation using spike timing dependent plasticity. *Neurocomputing* 330, 39–47. doi: 10.1016/j.neucom.2018.11.014
- Tavanaei, A., and Maida, A. S. (2015). A Minimal spiking neural network to rapidly train and classify handwritten digits in binary and 10-digit tasks. *Int. J. Adv. Res. Artif. Intellig.* 4:7. doi: 10.14569/IJARAI.2015.040701
- Tavanaei, A., and Maida, A. S. (2017). “Multi-layer unsupervised learning in a spiking convolutional neural network,” in *International Joint Conference on Neural Networks* (IEEE), 2023–2030.
- Thiele, J. C., Bichler, O., and Dupret, A. (2018). Event-based, timescale invariant unsupervised online deep learning with STDP. *Front. Comput. Neurosci.* 12:46. doi: 10.3389/fncom.2018.00046
- Thorpe, S., Delorme, A., and Van Rullen, R. (2001). Spike-based strategies for rapid processing. *Neural Netw.* 14, 715–725. doi: 10.1016/S0893-6080(01)00083-1
- Xiao, H., Rasul, K., and Vollgraf, R. (2017). Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms. *arXiv* [preprint]arXiv:1708.07747. doi: 10.48550/arXiv.1708.07747
- Xu, W., Wang, J., and Yan, X. (2021). Advances in memristor-based neural networks. *Front. Nanotechnol.* 3:645995. doi: 10.3389/fnano.2021.645995
- Yao, P., Wu, H., Gao, B., Tang, J., Zhang, Q., Zhang, W., et al. (2020). Fully hardware-implemented memristor convolutional neural network. *Nature* 577, 641–646. doi: 10.1038/s41586-020-1942-4
- Zenke, F., and Ganguli, S. (2018). SuperSpike: supervised learning in multilayer spiking neural networks. *Neural Comput.* 30, 1514–1541. doi: 10.1162/neco\_a\_01086
- Zenke, F., and Neftci, E. O. (2021). Brain-inspired learning on neuromorphic substrates. *Proc. IEEE* 109, 935–950. doi: 10.1109/JPROC.2020.3045625
- Zhang, T., Jia, S., Cheng, X., and Xu, B. (2021). Tuning convolutional spiking neural network with biologically plausible reward propagation. *Trans. Neural Netw. Learn. Syst.* 33, 7621–7631. doi: 10.1109/TNNLS.2021.3085966
- Zhao, D., Zeng, Y., Zhang, T., Shi, M., and Zhao, F. (2020). GLSNN: a multi-layer spiking neural network based on global feedback alignment and local STDP plasticity. *Front. Comput. Neurosci.* 14:576841. doi: 10.3389/fncom.2020.576841
- Zou, X., Xu, S., Chen, X., Yan, L., and Han, Y. (2021). Breaking the von Neumann bottleneck: architecture-level processing-in-memory technology. *Sci. China Inform. Sci.* 64:160404. doi: 10.1007/s11432-020-3227-1