

Forward Direct Feedback Alignment for Online Gradient Estimates of Spiking Neural Networks

Florian Bacho¹ and Dominique Chu^{1*}

^{1*}CEMS, School of Computing, University of Kent, Kennedy Building,
Canterbury, CT2 7NF, Kent, United Kingdom.

*Corresponding author(s). E-mail(s): d.f.chu@kent.ac.uk;
Contributing authors: f.bacho@outlook.fr;

Abstract

There is an interest in finding energy efficient alternatives to current state of the art neural network training algorithms. Spiking neural networks are a promising approach, because they can be simulated energy efficiently on neuromorphic hardware platforms. However, these platforms come with limitations on the design of the training algorithm. Most importantly, backpropagation cannot be implemented on those. We propose a novel neuromorphic algorithm, the *Spiking Forward Direct Feedback Alignment* (SF DFA) algorithm, an adaption of *Forward Direct Feedback Alignment* to train SNNs. SF DFA estimates the weights between output and hidden neurons as feedback connections. The main contribution of this paper is to describe how exact local gradients of spikes can be computed in an online manner while taking into account the intra-neuron dependencies between post-synaptic spikes and derive a dynamical system for neuromorphic hardware compatibility. We compare the SF DFA algorithm with a number of competitor algorithms and show that the proposed algorithm achieves higher performance and convergence rates.

Keywords: Spiking Neural Networks, Tradeoffs, Time-To-First-Spike, Error Backpropagation

1 Introduction

The backpropagation algorithm (BP) ([Rumelhart et al, 1986](#)) underpins a good part of modern neural network (NN) based AI. BP-based training algorithms continue to be the state of the art in many areas of machine learning ranging from benchmark problems such as the MNIST dataset ([LeCun et al \(2010\)](#)) to the most recent transformer-based architectures ([Launay et al, 2020](#)). While its success is undeniable, BP has some disadvantages. The main one is that BP

is computationally expensive and relies on sequential processing of layers during both the forward and backward pass, limiting its scope for parallelisation. This is sometimes called *backward locking* (Huo et al, 2018).

Furthermore, under BP the update of any particular connection weight in a neural network requires global information about the entire network. This entails intense data processing needs (Crafton et al, 2019) and consequently high energy consumption (Crafton et al (2019); Han and Yoo (2019)). Most importantly, BP is also not compatible with neuromorphic hardware platforms (Nefcici et al, 2017a), such as Loihi (Davies et al, 2018) or SpiNNaker (Furber et al, 2014), because neuromorphic hardware cannot save the past states of neurons or reverse time.

In the light of this, there has been some recent interest in alternatives to BP that alleviate these issues. In the context of SNNs, random feedback learning, and particularly the *Direct Feedback Alignment* (DFA) algorithm is an interesting alternative to BP for online gradient approximation (Nefcici et al, 2017b; Bellec et al, 2020; Rostami et al, 2022; Lee et al, 2020; Shrestha et al, 2019, 2021). For example, the *Event-Driven Random Backpropagation* (eRBP) (Nefcici et al, 2017b) and the *e-prop* (Bellec et al, 2020; Rostami et al, 2022) algorithms use random feedback connections to send error signals to hidden neurons at every time steps, thus continuously accumulating gradients during the inference. Unlike BP, DFA algorithms can be run efficiently on neuromorphic hardware. However, these algorithms constantly project errors to hidden neurons which can lead to significant computation and energy consumption on hardware.

Alternatives such as the *Spike-Train Level Direct Feedback Alignment* (ST-DFA) (Lee et al, 2020) or the *Error-Modulated Spike-Timing-Dependent Plasticity* (EMSTDP) (Shrestha et al, 2019, 2021) algorithms accumulate local changes of weights during the inference and send a single error signal to hidden neurons through random direct feedback connections. These approaches are more energy-efficient because the error projection is only performed once, thus requiring fewer operations. They have also been successfully implemented on neuromorphic hardware and demonstrated promising performance and lower energy consumption compared to offline training on GPUs (Lee et al, 2020; Shrestha et al, 2021; Rostami et al, 2022).

However, they have significant shortcomings. The most important one is that the performance of the algorithms is far behind the state of the art, i.e. far behind BP. It is generally thought that DFA algorithms rely on the alignment between weight and feedback connections (Nøkland, 2016; Refinetti et al, 2021). In the context of rate-coded neural networks it has recently been shown that the strong bias introduced by the random feedback connections affects the convergence rate of DFA. A solution for this was proposed in the form of the FDFA algorithm (Bacho and Chu, 2024), which enhanced DFA by learning the feedback matrices using Forward-AD. This proved to be highly successful in the context of rate-coded (non-spiking) neurons.

The main contribution of this paper is the *Spiking Forward Direct Feedback Alignment* (SFDFFA) algorithm and adaptation of FDFA to SNNs. While the main idea of the two algorithms are the same, a number of conceptual modifications were necessary for SFDFFA. Firstly, Forward-AD is not suitable for SNNs, and we use instead, as an alternative, *graded spikes* to estimate the weight between output and hidden neurons. Secondly, since spikes are non-differentiable, new different approaches are necessary in order compute gradients. Here we

choose an approach based on (Göltz et al, 2021) which computes gradients of spike times of neurons. Unlike, (Göltz et al, 2021), however, we allow multiple spikes per neurons, which in turn required a novel approaches to compute local gradient (see (Bacho and Chu, 2023) for a similar approach). Thirdly, as we found, the approach we took leads to a singularity where the the gradient diverges; we propose an *ad-hoc* solution for this. Using extensive simulations on a variety of benchmark problems, we demonstrate that SFDDFA is effective, outperforming previous neuromorphic algorithms on a variety of tasks. However, it remains consistently below the current state of the art, as set by BP.

2 Neuron model and notation

2.1 Leaky Integrate-and-Fire Neuron

In this work, we use networks of Leaky Integrate-and-Fire (LIF) neurons with the following membrane potential $u_i(t)$ dynamics for the post-synaptic neuron i :

$$u_i(t) = \sum_{j \in \mathcal{P}_i} w_{i,j} \sum_{t_j^z \in \mathcal{T}_j} \Theta(t - t_j^z) \left[\exp\left(\frac{t_j^z - t}{\tau}\right) - \exp\left(\frac{t_j^z - t}{\tau_s}\right) \right] - \vartheta \sum_{t_i^k \in \mathcal{T}_i} \Theta(t - t_i^k) \exp\left(\frac{t_i^k - t}{\tau}\right) \quad (1)$$

where t is the current time, \mathcal{P}_i is the set of pre-synaptic neurons j connected to the post-synaptic neuron i , t_j^k is the time of the k -th (postsynaptic) spike of neuron j , and $\mathcal{T}_j := \{t_j^1, t_j^2, \dots, t_j^{n_j}\}$ is the set of timings of all n_j postsynaptic spikes of neuron j , $w_{i,j} \in \mathbb{R}$ is the strength of the synaptic connection between the post-synaptic neuron i and the pre-synaptic neuron j , τ_s and τ are respectively the synaptic and membrane time constants, $\vartheta \in \mathbb{R}$ is the spiking threshold of the neuron and

$$\Theta(x) = \begin{cases} 1 & \text{if } x \geq 0 \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

is the Heavyside step function that only triggers events that occurred in the past.

In essence, the LIF neuron integrates pre-synaptic spikes over time through its synapses, thus affecting its membrane potential. When the membrane potential $u_i(t)$ reaches the neuron's threshold $\vartheta \in \mathbb{R}$ at time t , a post-synaptic spike at time is fired at the output of the neuron and the membrane potential is reset to zero.

2.2 Exact Gradients of SNNs using Error Backpropagation

Inspired by its success with DNNs, the error Backpropagation (BP) algorithm has been successfully adapted to SNNs using various approximations that sidestep the non-differentiability aspect and the lack of closed form solution of spikes (Bohté et al, 2000; Shrestha and Orchard, 2018; Zheng et al, 2020; Wu et al, 2018; Jin et al, 2018). More recent approaches have successfully derived exact gradients by using implicit functions describing

the relation between membrane potential and spike timings (Wunderlich and Pehle, 2021). Alternatively, more recent methods such as Fast & Deep (Göltz et al, 2021) and its multi-spike extension introduce constraints to isolate explicit and differentiable closed-form solutions for spike times (Mostafa, 2016; Comsa et al, 2020; Göltz et al, 2021). By fixing the decay parameters of the membrane potential in eq. 1 to $\tau = 2\tau_s$, a closed form solution for the spike timings t_i^k can be found.

$$t_i^x = \tau \ln \left(\frac{2a_i^x}{b_i^x + \chi_i^x} \right) \quad (3)$$

where

$$\begin{aligned} a_i^x &:= \sum_{j \in \mathcal{P}_i} w_{ij} \sum_{z=1}^x \exp \left(\frac{t_j^z}{\tau_s} \right) \\ b_i^x &:= \sum_{j \in \mathcal{P}_i} w_{ij} \sum_{z=1}^x \exp \left(\frac{t_j^z}{\tau} \right) - \vartheta \sum_{z=1}^x \exp \left(\frac{t_i^z}{\tau} \right) \end{aligned}$$

and

$$\chi_i^x := \sqrt{(b_i^x)^2 - 4a_i^x \vartheta} \quad (4)$$

This explicit expression of post-synaptic spike timings thus allows for exact differentiation — see (Göltz et al, 2021) for more details about this method.

In the context of event-based SNNs, the gradient of given a differentiable loss function \mathcal{L} of a neural network with L layers is defined by all the partial derivatives $\frac{\partial \mathcal{L}}{\partial w_{pq}}$, such as:

$$\frac{\partial \mathcal{L}}{\partial w_{pq}} = \partial_i \mathcal{L}^x \frac{\partial t_i^{x,(L)}}{\partial t_j^{y,(L-1)}} \frac{\partial t_j^{y,(L-1)}}{\dots} \dots \frac{\partial t_k^{z,(l)}}{\partial t_k^z} \frac{\partial t_k^z}{\partial w_{pq}^{(l)}} \quad (5)$$

To simplify the notation, we introduced the shorthand $\partial_i \mathcal{L}^k := \frac{\partial \mathcal{L}}{\partial t_i^{k,(L)}}$ and used the convention that repeated indices are summed over. Note that the t_i^x superscripted index runs from 1 to n_i . To simplify notation, we will henceforth omit the layer super-script whenever this can be inferred from context. We will henceforth refer to $\frac{\partial t_i^k}{\partial w_{pq}}$ as the *local derivative* of the spike timing.

Under BP, the update of weights $\Delta^{\text{BP}} w_{pq}$ is then proportional to $\frac{\partial \mathcal{L}}{\partial w_{pq}}$, with the learning rate η as the proportionality factor. While this version of BP has yielded good results on SNNs in a number of contexts (), it has two major drawbacks. (i) Updates using BP require information that is non-local to the neuron, which is difficult to achieve on neuromorphic hardware. (ii) BP requires the hardware to retain a memory of the past states of neurons, which is not compatible with neuromorphic principles.

A simple solution to this is the above mentioned DFA algorithm, whereby the updates of the weights are not proportional to the gradient of the loss function but are instead made according to the modified DFA learning rule

$$\Delta^{\text{DFA}} w_{pq} := \eta 1_x \partial_i \mathcal{L}^x b_{ik}^{(l)} \frac{\partial t_k^{z,(l)}}{\partial w_{pq}^{(l)}} 1_z \quad (6)$$

Here $b_{ik}^{(l)}$ are the elements of a randomly drawn, but fixed *feedback matrix* $\mathbf{B}^{(l)}$ for layer l . The symbol 1_x is an indexed constant of value 1 which implements the sum over the derivatives of the individual spikes.

DFA sidesteps the need for backpropagating errors backward through space and time and allows for online error computation. In addition, because all post-synaptic spikes receive the same projected error, a local gradient of spikes can be locally computed during the inference without immediately requiring an error signal. For this reason, DFA has been successfully implemented in several neuromorphic hardware (Lee et al, 2020; Shrestha et al, 2021; Rostami et al, 2022).

While efficient to compute, the performance of DFA is typically not competitive with BP. Previously (Bacho and Chu, 2024), we showed in the context of rate-coded neural networks that the performance can be increased substantially by updating the feedback matrices along with the weights themselves. In order to compute the update to the feedback matrix, we require the neuron to store an additional value (the *grade*, denoted by d_i) locally. The grade is initialised to zero, and then updated as follows:

- Upon receiving an input spike from neuron j , the grade is increased by $w_{ij_0} \cdot d_{j_0}$ (note the summation convention does not apply here).
- Upon generating the first spike at time t_i^0 a random number is drawn from a normal distribution and stored as p_i^0 .
- Upon generating a spike the grade is reset to zero.

During training, the feedback matrices are then updated as follows

As in FDFFA, we define an update rule for the feedback connection $b_{o,i}$ that performs an exponential average of the weight estimate, such as:

$$b_{ji} \leftarrow (1 - \alpha) b_{ji} + \alpha d_j p_i^1 \quad (7)$$

where $0 < \alpha < 1$ is the feedback learning rate.

In the case of a single hidden layer this learning rule has a clear interpretation. It is straightforward to show that the feedback matrices lead become over time similar to the forward weights w_{ij} .

$$\begin{aligned} d_o &= \sum_{t_o^k \in \mathcal{T}_o} d_o^k \\ &= \sum_{i \in \mathcal{P}_o} w_{o,i} \sum_{t_i^z \in \mathcal{T}_i} d_i^z \\ &= \sum_{i \in \mathcal{P}_o} w_{o,i} \sum_{t_i^k \in \mathcal{T}_i} p_i^k \\ &= \sum_{i \in \mathcal{P}_o} w_{o,i} p_i^1 \end{aligned} \quad (8)$$

Multiplying d_o by the perturbation p_i^1 then leads to an unbiased estimate of the weight $w_{o,i}$.

$$\begin{aligned}
\mathbb{E} [d_o p_i^1] &= \mathbb{E} \left[\sum_{j \in \mathcal{P}_o} w_{o,j} p_j^1 p_i^1 \right] \\
&= \sum_{j \in \mathcal{P}_o} w_{o,j} \mathbb{E} [p_j^1] \mathbb{E} [p_i^1] \\
&= w_{o,i}
\end{aligned} \tag{9}$$

In that sense, SFDFFA then approximates BP.

Algorithm 1 Inference of a single hidden neuron i with the Spiking Forward Direct Feedback Alignment algorithm.

- 1: **Input:** The neuron index i , the neuron weights \mathbf{w} , the reset current $\vartheta = \vartheta - u_{\text{rest}}$, the random perturbation $p \sim \mathcal{N}(0,1)$ and the set \mathcal{T}_{pre} of pre-synaptic graded spikes $(j, t_j^z, d_j^z) \in \mathcal{T}_{\text{pre}}$ sorted in time where j is a pre-synaptic neuron index, t_j^z is a timing and d_j^z is a spike grade.
 - 2: **Initialize:** the factors $a = 0$ and $b = 0$, the spike count k , the local gradient $\nabla_{\text{loc}} \mathbf{w} = \mathbf{0}$, the set of output events $\mathcal{T}_{\text{post}} = \{\}$, the internal directional derivatives $s = 0$, and the local derivatives synaptic traces $\mathbf{f} = \mathbf{0}$ and $\mathbf{h} = \mathbf{0}$.
 - 3: **for all** (j, t_j^z, d_j^z) **in** \mathcal{T}_{pre} **do**
 - 4: $f_j \leftarrow f_j \exp\left(\frac{t_j^z}{\tau_s}\right)$ {Update local derivative traces}
 - 5: $h_j \leftarrow h_j \exp\left(\frac{t_j^z}{\tau}\right)$
 - 6: $a \leftarrow a + w_j \exp\left(\frac{t_j^z}{\tau_s}\right)$ {Integrate pre-synaptic spike into quadratic factors}
 - 7: $b \leftarrow b + w_j \exp\left(\frac{t_j^z}{\tau}\right)$
 - 8: $s \leftarrow s + w_j d_j^z$ {Update internal directional derivative}
 - 9: **if** the neuron fires a valid spike **then**
 - 10: **continue**
 - 11: **end if**
 - 12: $t \leftarrow \tau \ln\left(\frac{2a}{b + \sqrt{b^2 - 4a\vartheta}}\right)$ {Compute post-synaptic spike timing}
 - 13: $\mathbf{t}' \leftarrow -\frac{\tau}{a_i^k \exp\left(\frac{t}{\tau_s}\right)} \left[\mathbf{h} \exp\left(\frac{-t}{\tau}\right) - \mathbf{f} \exp\left(\frac{-t}{\tau_s}\right) \right]$ {Compute spike derivative}
 - 14: $\nabla_{\text{loc}} \mathbf{w} \leftarrow \nabla_{\text{loc}} \mathbf{w} + \mathbf{t}'$ {Update local gradient}
 - 15: $\mathbf{h} \leftarrow \mathbf{h} - \frac{\vartheta}{\tau} \exp\left(\frac{t}{\tau}\right) \mathbf{t}'$ {Apply recurrence}
 - 16: $b \leftarrow b - \vartheta \exp\left(\frac{t}{\tau}\right)$ {Reset of membrane potential}
 - 17: $d \leftarrow s + p$ **if** $k = 0$ **else** s {Compute directional derivative as spike grade}
 - 18: add (i, t, d) to $\mathcal{T}_{\text{post}}$
 - 19: **end for**
 - 20: **return** $\mathcal{T}_{\text{post}}, \nabla_{\text{loc}} \mathbf{w}$
-

Algorithm 2 Weights update of a single hidden neuron i with the Spiking Forward Direct Feedback Alignment algorithm.

- 1: **Input:** the neuron weights \mathbf{w} , the local gradient $\nabla_{\text{loc}} \mathbf{w}$, the feedback connections \mathbf{b} , the random perturbation p , the output directional derivatives \mathbf{d} , the output errors δ and the learning rates λ and α .
 - 2: $\mathbf{w} \leftarrow \mathbf{w} - \lambda \delta \mathbf{b} \nabla_{\text{loc}} \mathbf{w}$ {Update weights}
 - 3: $\mathbf{b} \leftarrow (1 - \alpha) \mathbf{b} + \alpha p \mathbf{d}$ {Update feedback connections}
 - 4: **return** \mathbf{w}, \mathbf{b}
-

2.3 Computing the DFA update rule

In this section, we now describe how we implemented the SFDA learning rule in practice in the context of this paper. Many definitions exist for these local gradients of spikes. For example, the gradient of the membrane potential at spike times (Lee et al, 2020) or surrogate local gradients (Bellec et al, 2020) can be used. Other works, such as EMSTDP (Shrestha et al, 2021), use biologically-plausible local gradients like STDP or simple accumulations of the number of pre-synaptic spikes as in eRBP (Nefcici et al, 2017b). Here we show that within our framework we can explicitly compute the errors $\partial_i \mathcal{L}^x$ and the of the local gradient $\frac{\partial t_k^{z,(l)}}{\partial w_{pq}^{(l)}}$ in eq. 6. In the following two subsections we will show how to do this.

2.3.1 Computing the error

In the particular examples that we consider here, we will use a spike count n to decode the output, with a loss function defined as:

$$\mathbf{L}(n) := \sum_i (y_i - n_i)^2 \quad (10)$$

Here, y_i is a particular training example and n_i the actual spike count of the output neuron of the spiking network; the sum runs over a mini-batch.

The spike count n is a complicated function of the spike timings of the output neurons, and consequently the spike timings of hidden neurons as well, as per eq. 6. We can then see that for this particular choice the error term in eq. 6 can be split up as follows.

$$1_x \partial_i \mathcal{L}^x = \frac{\partial \mathcal{L}}{\partial n} 1_x \frac{\partial n}{\partial t_i^x} \quad (11)$$

Computing the partial derivatives with respect to spike timings is computationally expensive. To simplify the computation, we absorb those into the random feedback matrices, which gives the following simplified DFA update rule:

$$\Delta^{\text{DFA}} w_{pq} := \eta \frac{\partial \mathcal{L}}{\partial n} b_k^{(l)} \frac{\partial t_k^{z,(l)}}{\partial w_{pq}^{(l)}} 1_z \quad (12)$$

Here, again, we used the notation that repeated indices are summed over.

2.3.2 Computing the local gradient

Next, we describe how to compute the local gradient in eq. 6. We observe that in eq. 3 t_i^x depends explicitly on w_{ij} , and also depends on t_i^{x-1} , which in turn depends on w_{ij} and its preceding spikes (cf. eq. 4). In order to capture the gradient accurately, we therefore need to replace the partial derivative in eq. 6 by the total derivative of the spike timing with respect to the weight w_{pq} . To compute this, we use the general formula for the total derivative as an *ansatz*.

$$\frac{dt_i^x}{dw_{ij}} = \frac{\partial t_i^x}{\partial w_{ij}} + \frac{\partial t_i^x}{\partial t_i^{x-1}} \frac{dt_i^{x-1}}{dw_{ij}} + \dots + \frac{\partial t_i^x}{\partial t_i^1} \frac{dt_i^1}{dw_{ij}} \quad (13)$$

Substituting in the explicit expressions for t_i^x as given by eq. 3 and evaluating the derivatives, we find after some straightforward, albeit tedious algebra an expression for the total derivative of the spike timing,

$$\frac{dt_i^x}{dw_{ij}} = \frac{\tau}{a_i^x} \left[1 + \frac{\vartheta}{\chi_i^x} \exp\left(\frac{t_i^x}{\tau}\right) \right] f_{ij}^x - \frac{\tau}{\chi_i^x} h_{ij}^x, \quad (14)$$

where we used the shorthands

$$f_{ij}^x = \sum_{\{z:t_j^z \leq t_i^x\}} \exp\left(\frac{2t_j^z}{\tau}\right) \quad (15)$$

and

$$h_{ij}^x = \sum_{\{z:t_j^z \leq t_i^x\}} \exp\left(\frac{t_j^z}{\tau}\right) - \sum_{z=1}^x \left(\frac{\vartheta}{\tau} \exp\left(\frac{t_i^z}{\tau}\right) \frac{dt_i^z}{dw_{ij}} \right) \quad (16)$$

and used the condition that $\tau = 2\tau_s$.

In appendix B, we derive a mathematically equivalent, but computationally more convenient and faster form of this equation.

$$\frac{dt_i^x}{dw_{ij}} = \frac{\tau}{a_i^x \exp\left(\frac{-2t_i^x}{\tau}\right) - \vartheta} \left[f_{ij}^x \exp\left(\frac{-2t_i^x}{\tau}\right) - h_{ij}^x \exp\left(\frac{-t_i^x}{\tau}\right) \right] \quad (17)$$

2.4 Critical Points of Local Gradients

It is apparent from eq. 14 that the local gradient diverges when $\chi \rightarrow 0$ or equivalently when $a_i^x \exp\left(\frac{-t_i^x}{\tau_s}\right) = \vartheta$ (cf. (Wunderlich and Pehle, 2021; Takase et al, 2009a,b) for similar observation). A common solution to mitigate gradient explosion is adding additional mechanisms to limit the size of the gradient such as *gradient clipping* (Wunderlich and Pehle, 2021; Hong et al, 2020). However, gradient clipping requires computing the norm of the gradient, which cannot be locally performed on neuromorphic hardware. In this work, we propose a simple *ad-hoc* modification of the local gradient to remove the cause of the critical points.

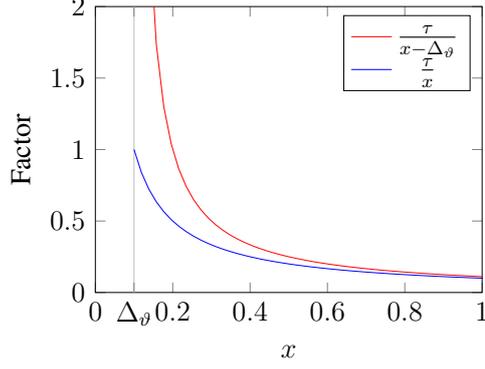


Fig. 1: Original factor $\frac{\tau}{x - \Delta\vartheta}$ (red line) and the modified factor $\frac{\tau}{x}$ as a function of $x = a_i^k \exp\left(\frac{-t_i^k}{\tau_s}\right)$. As $a_i^k \exp\left(\frac{-t_i^k}{\tau_s}\right)$ approaches ϑ , the original factor diverges towards infinity while the modified factor is bounded.

Our solution is to drop ϑ from eq. 17, thus obtaining:

$$\frac{dt_i^x}{dw_{ij}} = \frac{\tau}{a_i^x \exp\left(\frac{-t_i^x}{\tau_s}\right)} \left[f_{ij}^x \exp\left(\frac{-t_i^x}{\tau_s}\right) - h_{ij}^x \exp\left(\frac{-t_i^x}{\tau}\right) \right] \quad (18)$$

Here, $\frac{\tau}{a_i^x \exp\left(\frac{-t_i^x}{\tau_s}\right)} \leq \frac{\tau}{\vartheta}$ has an upper bound and does not diverge towards infinity (see Figure 1). While this creates an additional bias in the gradient estimates, dropping the term ϑ avoids exploding gradients and improves the stability of the weights updates.

2.5 Neuromorphic Hardware Compatibility

The local gradient in Equation 18 is adapted for computation on von Neumann computers. Therefore, to make the SFDEFA algorithm compatible with neuromorphic hardware, we derive an eligibility trace that can be evaluated with a system of ordinary differential equations on hardware.

Eq. 18 can be written as:

$$\frac{dt_i^x}{dw_{ij}} = -\frac{\tau}{I_i(t_i^x)} e_{ij}(t_i^x) \quad (19)$$

where

$$\begin{aligned} I_i(t) &= a_i^x \exp\left(\frac{-t}{\tau_s}\right) \\ &= \sum_{j \in \mathcal{P}_i} w_{ij} \sum_{\{z: t_j^z \leq t\}} \alpha(t - t_j^z) \end{aligned} \quad (20)$$

is the synaptic current at the spike time, which is computed by the LIF neuron (Gerstner and Kistler, 2002), and

$$\begin{aligned}
e_{ij}(t) = & \sum_{\{z:t_j^z \leq t\}} \underbrace{\left[\exp\left(\frac{t_j^z - t}{\tau}\right) - \exp\left(\frac{t_j^z - t}{\tau_s}\right) \right]}_{=\epsilon(t-t_j^z)} \\
& - \sum_{\{z:t_j^z \leq t\}} \vartheta \underbrace{\exp\left(\frac{t_i^z - t}{\tau}\right)}_{=\eta(t-t_i^z)} \frac{dt_i^x}{dw_{ij}} \frac{1}{\tau}
\end{aligned} \tag{21}$$

is an eligibility trace computed over time that coincides with the desired expression required by the local gradient at the spike time t_i^k . We can observe that Equation 21 has a form that is similar to the mapping of the LIF neuron to the SRM model. Therefore, this expression can be implemented on neuromorphic hardware using a LIF model at each synapse and locally be used at post-synaptic spike times for the computation of local gradients.

3 Empirical Results

We now present our empirical results with our proposed SF DFA algorithm. We empirically highlight the critical points in the exact local gradient of spikes and demonstrate the effectiveness of our modified local gradient in preventing gradient explosions. We then compare the performance and convergence rate of SF DFA with DFA and BP on several benchmark datasets as well as the weight and gradient alignment of DFA and SF DFA.

3.1 Critical Points Analysis

We now analyze the stability of the exact local gradient of spikes and compare it with our modified local gradient, defined in Equation 18. To gain an intuitive understanding of the problem, we first consider a toy example of a two-input neuron receiving four pre-synaptic spikes. We evaluate the local gradient with the exact derivative of spikes (Equation 17) and with our modified local gradient (Equation 18). Figure 2 shows the respective gradient fields computed using the exact and modified local gradients. It can be observed in Figure 2a that the exact local gradient contains several critical points where its norm is significantly larger than neighboring vectors.

To investigate the cause behind these large norms, we examined the neuron’s internal state over time at one of these critical points (marked in red in Figure 2). Figure 3 depicts the temporal evolution of the membrane potential, input current, factor $\frac{\tau}{I_i(t) - \vartheta}$, and the computed local gradient. Notably, it can be observed that the last post-synaptic spike fired by the neuron narrowly crosses the threshold (i.e. hair trigger) due to a low input current. This low current leads to the factor $\frac{\tau}{I_i(t) - \vartheta}$ becoming large due to its divergence when $I_i(t)$ approaches ϑ (see Figure 1). Consequently, the local gradient at this spike time significantly increases.

On the other hand, Figure 2 demonstrates that our proposed modified local gradient maintains a consistent norm across the weight space without any instability points where the norm deviates abnormally from neighboring vectors. Additionally, Figure 3 illustrates that, at the

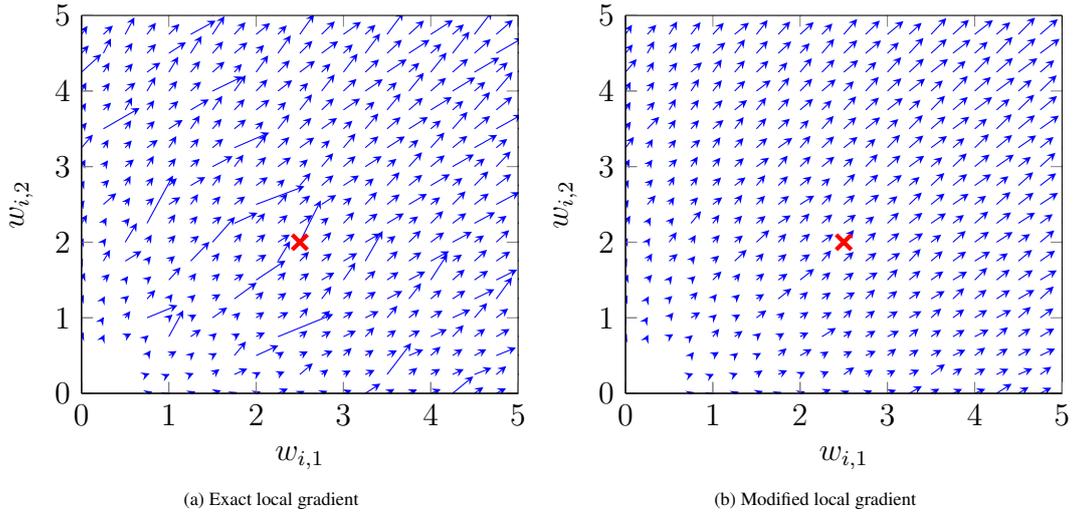


Fig. 2: Gradient fields of a single neuron with two inputs computed using the exact local gradients (a) and the modified local gradients (b). We can observe that the exact local gradient contains critical points where the norm is abnormally large compared to neighboring regions. However, with the modified local gradients, the norm of the gradient is consistent throughout the weight space, mitigating the gradient explosion caused by the critical points. The red crosses at correspond to the critical point visualized in Figure 3. See main text for details on how this figure was generated. For these simulations, we used $\tau_s = 0.01$, $\theta = 0.01$ and the time window was of length 0.01 seconds.

spike time when the membrane potential narrowly reaches the threshold, the modified factor $\frac{\tau}{I_i(t)}$ exhibits a significantly lower value compared to the original factor $\frac{\tau}{I_i(t) - \bar{v}}$. This is because the modified factor is upper-bounded, as shown in Figure 1. As a result, the contribution of this spike to the modified local gradient is substantially reduced in comparison to its contribution to the exact local gradient.

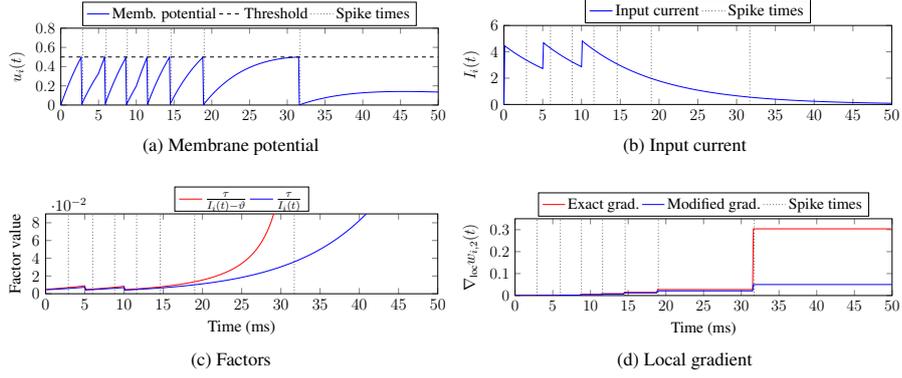


Fig. 3: Evolution over time of the membrane potential (a), input current (b), gradient factors (c) and local gradients (d) at an instability point ($w_{i,1} = 2.5$ and $w_{i,2} = 2.0$ in Figure 2) in a minimal example of a single spiking neuron with two inputs. Here all vertical dotted lines correspond to post-synaptic spike times. The red lines in both (c) and (d) represent the original factor and local gradient, while the blue lines represent the modified factor and local gradient. These figures show that the last post-synaptic spike fired by the neuron occurs when the membrane potential narrowly reaches the threshold. This narrow threshold crossing is attributed to the low input current $I_i(t)$ which causes large factors $\frac{\tau}{I_i(t) - \theta}$ and consequently large local gradients. In contrast, our modified factor $\frac{\tau}{I_i(t)}$ restricts the amplitude of the factor, thereby moderating the scale of the gradient. For these simulations, we used $\tau_s = 0.01$, $\theta = 0.01$ and the time window was of length 0.01 seconds.

To assess the impact of critical points on the convergence of SNNs, we conducted an experiment using the MNIST dataset. Identical SNNs were trained, both containing two layers of fully connected neurons and sharing the same hyperparameters. These networks were initialized with identical weights, differing only in the type of local gradient used for training. One network was trained with the exact local gradient, while the second network was trained with the modified gradient defined in Equation 18.

Figure 4 shows the evolution of the training loss, training accuracy, test loss, and test accuracy for each network during the training of each network. It can notably be observed that the network trained with the modified local gradient converges slightly faster than the SNN trained with the exact local gradient. This suggests that the modified gradient mitigates the impact of instabilities on convergence, leading to enhanced learning.

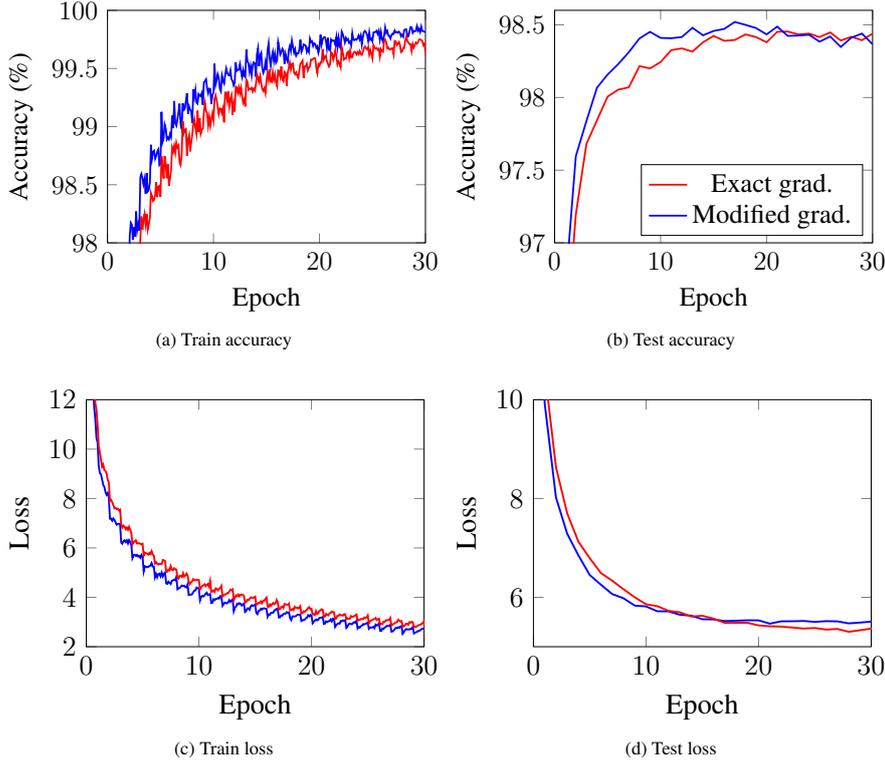


Fig. 4: Evolution of the training accuracy (Figure 4a), test accuracy (Figure 4b), train loss (Figure 4c) and test loss (Figure 4d) during the training of a two-layer SNN on the MNIST dataset. Red lines correspond to the metrics of the SNN updated using the exact local gradient while blue lines correspond to the metrics of the SNN updated using the modified local gradient defined in Equation 18. These figures show that the modified gradient converges slightly faster than the exact local gradient.

3.2 Performance

We evaluate the performance of the proposed SF DFA algorithm as well as BP and DFA in SNNs with four benchmark datasets, namely MNIST (LeCun et al, 2010), EMNIST (Cohen et al, 2017), Fashion MNIST (Xiao et al, 2017) and the *Spiking Heidelberg Digits* (SHD) dataset (Cramer et al, 2022). To encode the pixel values of static images (e.g. MNIST, EMNIST, and Fashion MNIST) into spikes that can be used as inputs of SNNs, we applied a coding scheme with an encoding window of $T = 100$ ms. To do this, we normalised the pixel values and calculated the spike time t for a pixel value x as

$$t = T(1 - x) \quad (22)$$

where T is the length of encoding window, which was chosen to be 100ms throughout. For pixel values of 0 (i.e. black), no spikes were generated. In this way, TTFS therefore produces sparse temporal encoding of greyscale images that are fast to process in my event-based simulator. For the SHD dataset, I used the spike trains as provided by the dataset.

We used a spike count strategy to decode the output spike counts of the SNNs. Specifically, each output neuron of the network encoded a possible category for the data. The target spike count for the correct category was set to 10, whereas the target spike count for the incorrect category was set to 3. Note that we did not choose a target of 0 for the latter to avoid dead neurons. We then used a spike count *Mean Squared Error* loss function for training. We benchmarked our method with fully-connected SNNs of different sizes. For MNIST and EMNIST, we used two-layer SNNs with 800 hidden neurons, for Fashion MNIST, a three-layer SNN with 400 hidden neurons per hidden layer, and a two-layer SNN with 128 hidden neurons for the SHD dataset. We also compared it with BP applied to an SNN, as described in [Bacho and Chu \(2023\)](#),

Table 1 summarizes the performance of each method on each dataset. We can observe that both the DFA and SF DFA algorithms achieve test performances similar to that of BP for the MNIST, EMNIST, and Fashion MNIST datasets. However, for the SHD dataset, neither algorithm attains a level of accuracy comparable to BP. This indicates that the direct feedback learning approach with a single error signal fails to effectively generalize when applied to temporal data.

Table 1: Average best test performance of BP, DFA and the proposed SF DFA on the MNIST, EMNIST, Fashion MNIST and SHD datasets.

Dataset	Architecture	BP	DFA	SF DFA
MNIST	800-10	98.88 \pm 0.02%	98.42 \pm 0.06%	98.56 \pm 0.04%
EMNIST	800-47	85.75 \pm 0.06%	79.48 \pm 0.11%	82.33 \pm 0.10%
Fashion MNIST	400-400-10	90.19 \pm 0.12%	89.41 \pm 0.12%	89.73 \pm 0.17%
SHD	128-20	66.79 \pm 0.66%	52.70 \pm 2.30%	54.63 \pm 1.16%

Despite this limitation, the proposed SF DFA algorithm consistently outperforms DFA across all benchmarked datasets. However, there remains a noticeable gap between the performance of SF DFA and BP, particularly as the complexity of the task increases.

3.3 Convergence

The F DFA algorithm introduced in ([Bacho and Chu, 2024](#)) demonstrated improvements in terms of convergence compared to DFA. To evaluate if this is also the case in SNNs, we recorded the train loss, train accuracy, test loss, and test accuracy of SNNs trained with the DFA and SF DFA algorithms on each dataset.

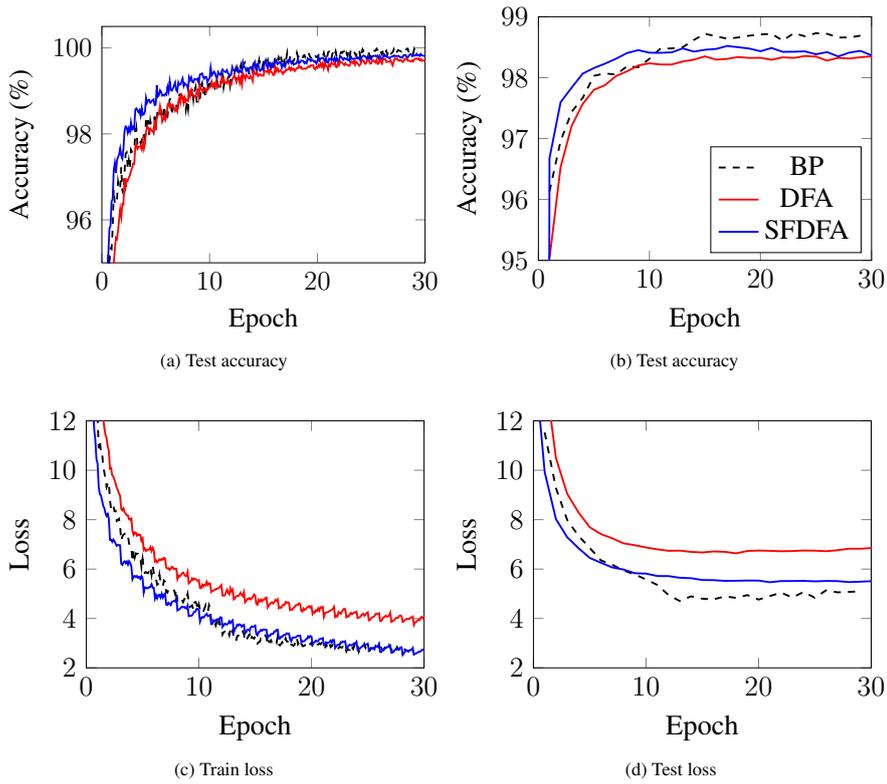


Fig. 5: Evolution of the averaged train accuracy (Figure 5a), test accuracy (Figure 5b), train loss (Figure 5c) and test loss (Figure 5d) during the training of two-layers SNNs on the MNIST dataset. Black dashed lines correspond to BP. Blue and red solid lines correspond to the SFDFA and DFA algorithms respectively.

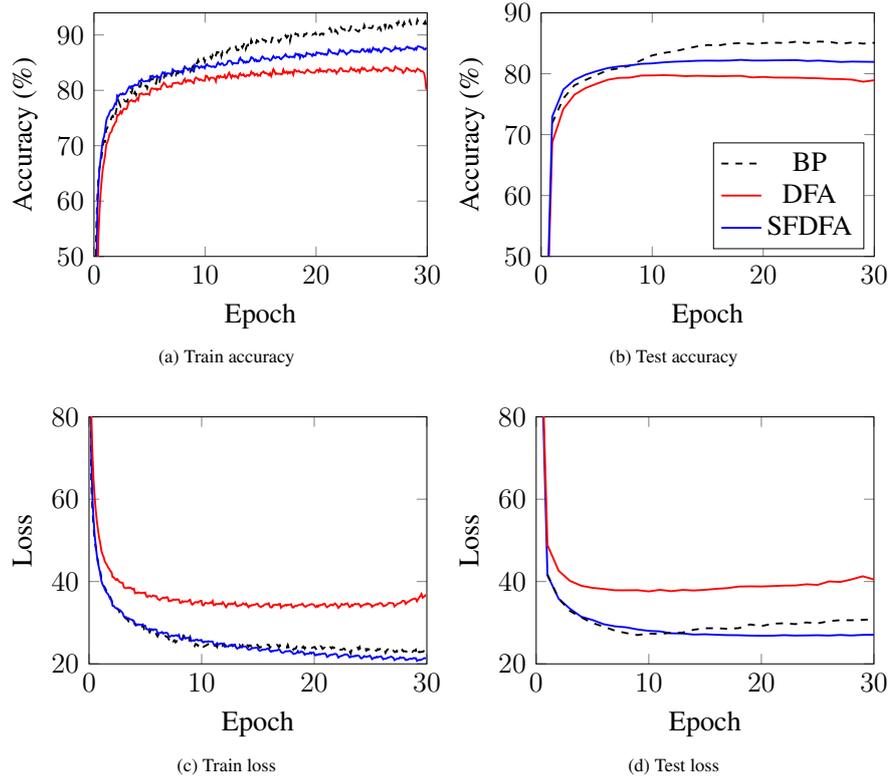


Fig. 6: Evolution of the averaged train accuracy (a), test accuracy (b), train loss (Figure 5c) and test loss (Figure 5d) during the training of two-layers SNNs on the EMNIST dataset. Black dashed lines correspond to BP. Blue and red solid lines correspond to the SF DFA and DFA algorithms respectively.

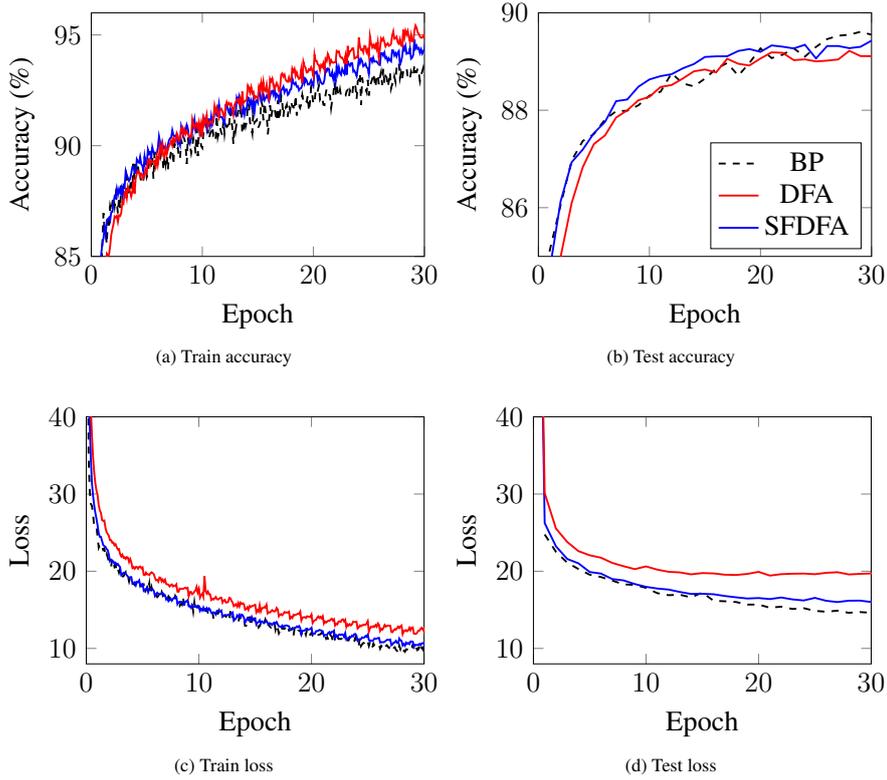


Fig. 7: Evolution of the averaged train accuracy (a), test accuracy (b), train loss (Figure 5c) and test loss (d) during the training of two-layers SNNs on the Fashion MNIST dataset. Black dashed lines correspond to BP. Blue and red solid lines correspond to the SFDFA and DFA algorithms respectively.

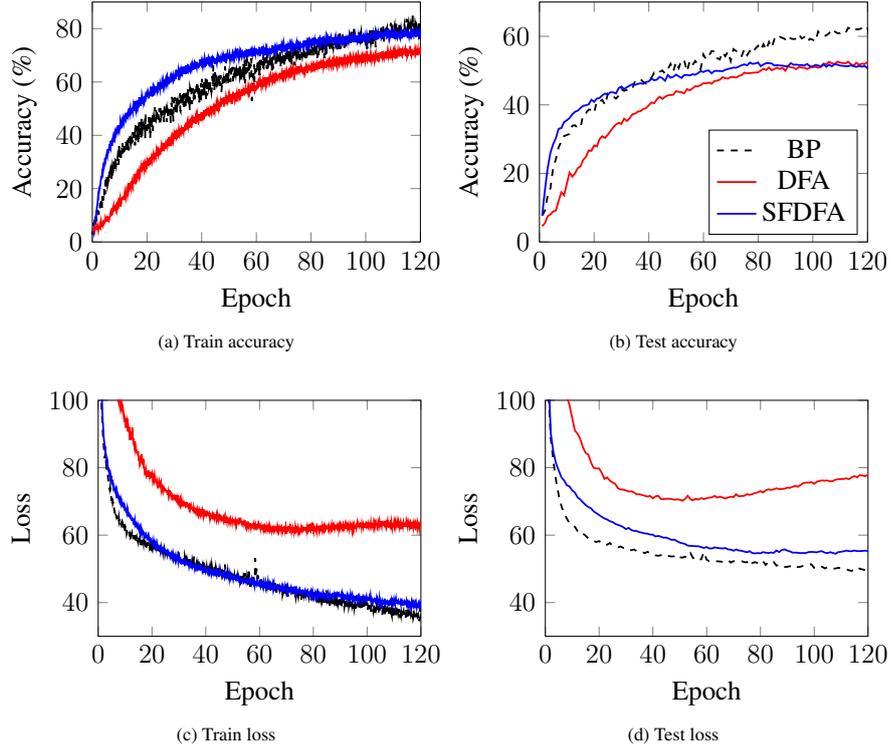


Fig. 8: Evolution of the averaged train accuracy (a), test accuracy (b), train loss (c) and test loss (d) during the training of two-layers SNNs on the Spiking Heidelberg Digits dataset. Black dashed lines correspond to BP. Blue and red solid lines correspond to the SFDDFA and DFA algorithms respectively.

Figures 5, 6, 7 and 8 show the evolution of the different metrics with the MNIST, EMNIST, Fashion MNIST and SHD datasets respectively. It can be observed that a gap exists in both the train and test loss between the BP and DFA algorithms. In contrast, the train and test loss for the proposed SFDDFA algorithm closely follow the loss values of BP, indicating similar convergence rates. Interestingly, during the initial stage of training, both the train and test accuracies of the SFDDFA algorithm increase faster than both BP and DFA. This behavior is particularly prominent in Figures 5 and 8. However, in the later stage of training, the test accuracy of SFDDFA plateaus while the test accuracy of BP continues to improve. This suggests that the proposed SFDDFA algorithm does not generalize as much as BP on test data. However, compared to DFA, the test accuracy of the SFDDFA algorithm increases significantly faster, especially on temporal data (see Figure 8). Altogether, we find that the SFDDFA algorithm shows better convergence than DFA, especially for more difficult tasks.

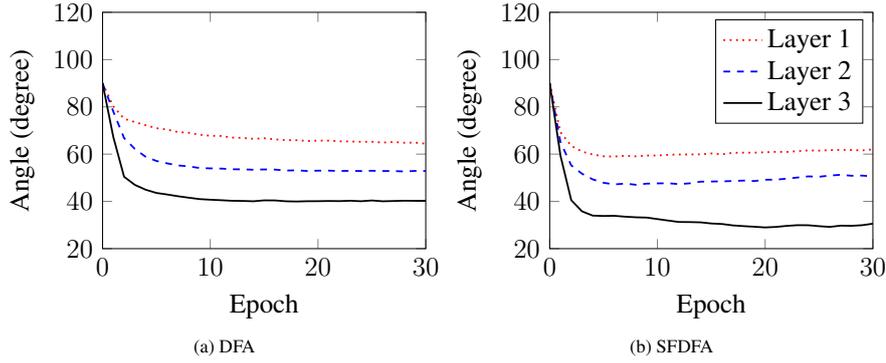


Fig. 9: Layerwise alignment between spiking gradient estimates and the true gradient computed using BP. These figures show that the SFDDFA algorithm Figure (b) aligns better with the true gradient than DFA (a), especially in layers close to the outputs.

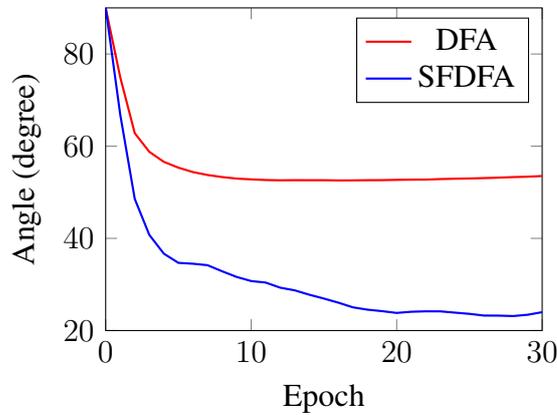


Fig. 10: Evolution of the alignment between the output weights and last hidden layer feedbacks in a 4-layer SNN trained on MNIST with DFA (red line) and SFDDFA (blue line). The weights and feedback connections align faster and better with the SFDDFA algorithm than with DFA. Moreover, the weights alignment in SNNs correlates with the gradient alignment (see Figure 9).

3.4 Weights and Gradient Alignment

We measured the bias of the gradient estimates by recording the layer-wise alignment between the approximate gradients and the true gradients computed using BP (see (Bacho and Chu, 2024)). For this experiment, we trained a 4-layer SNN on the MNIST dataset for 30 epochs using both the DFA and SFDDFA algorithms to compute gradient estimates. We then calculated the angles between these estimates and the true gradient computed by BP.

The evolution of the resulting angles during training is given in Figure 9. It can be observed that the gradient estimates provided by FDDFA align faster and exhibit a lower angle

with respect to BP when compared with DFA. This indicates that the proposed SF DFA algorithm achieves a lower level of bias earlier than DFA, with weight updates that better follow BP. However, it is important to note that the gap in alignment between SF DFA and DFA diminishes with the depth of the layer. Specifically, when comparing the alignments of the third layer (i.e. closest to the output) trained with SF DFA and DFA, the former demonstrates significantly better alignment. However, in the case of the first layer, both methods exhibit similar levels of alignment. This suggests that the benefits of SF DFA become more pronounced in layers close to the outputs.

In contrast with the FDFA algorithm which estimates derivatives, the SF DFA algorithm approximates the derivatives of spikes by estimating the weights connections between hidden and output neurons. Therefore, in addition to the layer-wise gradient alignment, we measured the angle between the vectors represented by the output weights and the vector represented by the feedback connections for the last hidden layer. Also known as *weight alignment*, it is believed to be the main source of gradient alignment in DFA (Refinetti et al, 2021). Figure 10 shows the evolution of the alignment between the flattened output weights and the flattened feedback connections to the last hidden layer. The alignments for deeper hidden layers were ignored as the exact forward weights for these layers are unknown. We can see in this figure that the output weight and the feedback connections trained with the SF DFA algorithm align faster and better than those trained with DFA. Moreover, we can observe that the weight alignment correlates with the gradient alignment of the third layer in Figure 9. This suggests that estimating the forward weights as feedback connections makes the approximate gradients align with the true gradients.

4 Discussion

In this paper, we proposed the SF DFA algorithm, a spiking adaptation of FDFA that trains SNNs in an online and local manner. The proposed algorithm computes local gradients of post-synaptic spikes by taking into account all intra-neuron dependencies and uses direct feedback connections to linearly project output errors to hidden neurons without unrolling the neuron’s dynamics through space and time. Similarly to the FDFA algorithm, SF DFA estimates the derivatives between hidden and output neurons as feedback connections by propagating directional derivatives as spike grades during the inference. More precisely, SF DFA estimates the weights between output and hidden layers and ignores the temporal relationships between spikes to avoid large variances that could hinder the convergence of feedback connections.

We also demonstrated the existence of critical points where the norm of the exact local gradient diverges towards infinity. These critical points have previously been discovered in other formulations of exact gradients (Takase et al, 2009a,b; Wunderlich and Pehle, 2021) of spikes. In our work, we identified the cause of these critical points in the computation of the exact local gradients and proposed a simple modification of the derivatives that suppresses gradient explosions. While this ad hoc solution introduces an additional bias to the gradient estimates, we showed that it enabled faster convergence of SNNs than the exact gradient when trained with SF DFA. This implies that the rate of convergence of our algorithm benefits from the improved stability of our modified local gradient.

Our empirical results showed that the proposed SF DFA consistently converges faster and achieves higher performance than DFA on all benchmark datasets. However, while our method also performs better than DFA on the SHD dataset, a significant gap still exists with BP. This suggests that the learning procedure used in both DFA and SF DFA has limitations when applied to highly temporal data. Therefore, future work could explore alternative update methods to improve this performance gap with BP on temporal data.

In our experiments, we measured the alignment between the approximate gradients computed by DFA and SF DFA and the true gradient computed by BP. Our results showed that the approximate gradients computed by the proposed SF DFA algorithm align faster with the true gradient than DFA. This suggests that weights are updated with steeper descending directions in SF DFA than in DFA. This could explain the increased convergence rate experienced by our algorithm. However, we observed that the gradient estimates in SNNs align less than the FDFA algorithm applied to DNNs. This weak alignment can be explained by two factors. First, by bounding the local derivatives, the modified local gradient of our method slightly changes the direction of the weight updates. Second, the complex temporal relationships between spikes are ignored in the computation of the directional derivatives to avoid introducing large variances in the feedback updates. Ignoring these temporal relationships could make the approximate gradients further deviate from the true gradient.

In addition to the gradient alignment, we measured the alignment between the network weights and feedback connections in both DFA and SF DFA. We observed a stronger difference in weight alignment between DFA and SF DFA than in gradient alignment. This could be explained by the fact that our method estimates the weight connections between output and hidden neurons as feedback rather than derivatives. In particular, we observed that the weight alignment of SF DFA correlates with the gradient alignment of the last layer, suggesting that estimating weights contributes to the gradient alignments. However, it is still unclear why deeper layers fail to align more than in DFA. Future work could therefore focus on improving the gradient alignment of deep layers to improve the rate of convergence as well as the performance of SNNs when trained with SF DFA.

From an engineering point of view, the local gradients of spikes can be locally computed by neurons by implementing dedicated circuits that evaluate the dynamical system of the LIF neuron. Moreover, the computation and propagation of directional derivatives during the inference can be implemented through the grades of spikes. Spike grades are features that have recently been added to several large-scale neuromorphic platforms such as Loihi 2 (Frady et al, 2022) and SpiNNaker (Furber et al, 2014). Traditionally used to modify the amplitude of spikes for computational purposes, our work instead proposes the use of spike grades for learning purposes. Finally, direct feedback connections have widely been implemented on various neuromorphic platforms, thus supporting the hardware compatibility of SF DFA.

Therefore, by successfully addressing the limitations of BP, the proposed SF DFA algorithm represents a promising step towards the implementation of neuromorphic gradient descent. While there are still areas for improvement and exploration, our findings contribute to the growing body of knowledge aimed at improving the field of neuromorphic computing.

Acknowledgements. This study was part funded by EPSRC grant EP/T008296/1.

Data and code availability. All data used in this paper is publicly available benchmark data and has been cited in the main text (LeCun et al, 2010; Xiao et al, 2017; Cramer et al, 2022).

Conflict of interest. The authors declare that there is no conflict of interest.

Appendix A Experimental Settings

In this section, we describe the experimental settings used to produce our empirical results, including the benchmark datasets, encoding and decoding, network architectures, the loss function, hyperparameters as well as the software and hardware settings.

A.1 Firing Rate Regularization

Without additional constraint, neurons may exhibit high firing rates to achieve lower loss values which could increase energy consumption and computational requirements. To prevent this issue, we implemented a firing rate regularization that drives the mean firing rate of neurons towards a given target during training. By incorporating firing rate regularization, the neural network is encouraged to find a balance between learning from the data and avoiding high firing rates. This can lead to improved generalization, reduced energy consumption, and enhanced stability.

Formally, a penalty term for high firing rates is added to the loss function, such as:

$$\mathcal{L}_{\text{reg}}(\mathbf{x}) = \mathcal{L}(\mathbf{x}) + \beta \sum_i \left(\mathbb{E}_{\mathbf{x}}[n_i] - \hat{n} \right)^2 \quad (\text{A1})$$

where $\beta > 0$ is a constant defining the strength of the regularization and $\mathbb{E}_{\mathbf{x}}[n_i]$ is the mean firing rate of the hidden neuron i . The mean firing rate can be estimated in an online manner by computing a moving average or an exponential average of the firing rate over the last inference. If batch learning is used, a mean firing rate can be computed from the batch.

A.2 Update Method and Hyperparameters

We used the Adam (Kingma and Ba, 2017) algorithm to update both the weights and the feedback connections. We used the default values of $\beta_1 = 0.9$, $\beta_2 = 0.999$ and $\epsilon = 10^{-8}$ (Kingma and Ba, 2017) and a batch size of 50 for fully-connected SNNs. We used a learning rate of $\lambda = 0.003$ for image classification and $\lambda = 0.001$ for audio classification with the SHD dataset. In addition, we used a feedback learning rate of $\alpha = 10^{-4}$ in every experiment.

Experimental conditions were standardized for BP, DFA and SF DFA. We used the same hyperparameters as the method proposed in

A.3 Event-Based Simulations on GPU

To simulate and train SNNs, we reused the event-based simulator used in (Bacho and Chu, 2023). We adapted the GPU kernels related to the inference to compute local gradients as well as directional derivatives in an online manner. Moreover, we replaced the code performing error backpropagation with feedback learning. Similarly to neuromorphic hardware, our simulator never backpropagates errors backward through time and performs all computations in an online manner.

Appendix B Deriving eq. 17

To reduce the computational requirements of the local gradient evaluation, Equation 14 can be further simplified by re-introducing the post-synaptic spike time t_i^k into the equation:

$$\begin{aligned} \frac{dt_i^x}{dw_{ij}} &= \frac{\exp\left(\frac{-t_i^x}{\tau_s}\right) \tau}{\exp\left(\frac{-t_i^x}{\tau_s}\right) a_i^x} \left[1 + \frac{\exp\left(\frac{-t_i^x}{\tau}\right) \vartheta}{\exp\left(\frac{-t_i^x}{\tau}\right) x_i^x} \exp\left(\frac{t_i^x}{\tau}\right) \right] f_i^x - \frac{\exp\left(\frac{-t_i^x}{\tau}\right) \tau}{\exp\left(\frac{-t_i^x}{\tau}\right) x_i^x} h_i^x \\ &= \frac{\tau}{a_i^x \exp\left(\frac{-t_i^x}{\tau_s}\right)} \left[1 + \frac{\vartheta}{x_i^x \exp\left(\frac{-t_i^x}{\tau}\right)} \right] f_i^x \exp\left(\frac{-t_i^x}{\tau_s}\right) - \frac{\tau}{x_i^x \exp\left(\frac{-t_i^x}{\tau}\right)} h_i^x \exp\left(\frac{-t_i^x}{\tau}\right) \end{aligned} \quad (\text{B2})$$

Then we can isolate an expression for $x_i^x \exp\left(\frac{-t_i^x}{\tau}\right)$, such as:

$$\begin{aligned} t_i^x &= \tau \ln\left(\frac{2a_i^x}{b_i^x + x_i^x}\right) \\ \Leftrightarrow \exp\left(\frac{t_i^x}{\tau}\right) &= \frac{2a_i^x}{b_i^x + x_i^x} \\ \Leftrightarrow \exp\left(\frac{t_i^x}{\tau}\right) \exp\left(\frac{-t_i^x}{\tau}\right) &= \exp\left(\frac{-t_i^x}{\tau}\right) \frac{\exp\left(\frac{-t_i^x}{\tau}\right) 2a_i^x}{\exp\left(\frac{-t_i^x}{\tau}\right) b_i^x + x_i^x} \\ \Leftrightarrow \exp\left(\frac{t_i^x - t_i^x}{\tau}\right) &= \frac{2a_i^x \exp\left(\frac{-t_i^x}{\tau_s}\right)}{(b_i^x + x_i^x) \exp\left(\frac{-t_i^x}{\tau}\right)} \\ \Leftrightarrow 1 &= \frac{2a_i^x \exp\left(\frac{-t_i^x}{\tau_s}\right)}{(b_i^x + x_i^x) \exp\left(\frac{-t_i^x}{\tau}\right)} \\ \Leftrightarrow x_i^x \exp\left(\frac{-t_i^x}{\tau}\right) &= a_i^x \exp\left(\frac{-t_i^x}{\tau_s}\right) - \vartheta \end{aligned} \quad (\text{B3})$$

as of the constraint between the synaptic and membrane time constant $\tau = 2\tau_s \Leftrightarrow \exp\left(\frac{-t_i^x}{\tau}\right)^2 = \exp\left(\frac{-t_i^x}{\tau_s}\right)$.

$$\begin{aligned} u(t_i^x) &= b_i^x \exp\left(\frac{-t_i^x}{\tau}\right) - a_i^x \exp\left(\frac{-t_i^x}{\tau_s}\right) = \vartheta \\ \Leftrightarrow b_i^x \exp\left(\frac{-t_i^x}{\tau}\right) &= a_i^x \exp\left(\frac{-t_i^x}{\tau_s}\right) + \vartheta \end{aligned} \quad (\text{B4})$$

Therefore, the local gradient $\frac{dt_i^x}{dw_{ij}}$ of spikes simplifies to:

$$\begin{aligned} \frac{dt_i^x}{dw_{ij}} &= \frac{\tau}{a_i^x \exp\left(\frac{-t_i^x}{\tau_s}\right)} \left[\frac{a_i^x \exp\left(\frac{-t_i^x}{\tau_s}\right) - \vartheta}{a_i^x \exp\left(\frac{-t_i^x}{\tau_s}\right) - \vartheta} + \frac{\vartheta}{a_i^x \exp\left(\frac{-t_i^x}{\tau_s}\right) - \vartheta} \right] f_i^x \exp\left(\frac{-t_i^x}{\tau_s}\right) \\ &\quad - \frac{\tau}{a_i^x \exp\left(\frac{-t_i^x}{\tau_s}\right) - \vartheta} h_i^x \exp\left(\frac{-t_i^x}{\tau}\right) \\ &= \frac{\tau}{a_i^x \exp\left(\frac{-t_i^x}{\tau_s}\right) - \vartheta} \left[f_i^x \exp\left(\frac{-t_i^x}{\tau_s}\right) - h_i^x \exp\left(\frac{-t_i^x}{\tau}\right) \right] \end{aligned} \quad (\text{B5})$$

References

- Bacho F, Chu D (2023) Exploring tradeoffs in spiking neural networks. *Neural Computation* 35(10):1627–1656. URL <https://kar.kent.ac.uk/101545/>
- Bacho F, Chu D (2024) Low-variance forward gradients using direct feedback alignment and momentum. *Neural Networks* 169:572–583. <https://doi.org/https://doi.org/10.1016/j.neunet.2023.10.051>, URL <https://www.sciencedirect.com/science/article/pii/S0893608023006172>
- Bellec G, Scherr F, Subramoney A, et al (2020) A solution to the learning dilemma for recurrent networks of spiking neurons. *Nature Communications* 11(1). <https://doi.org/10.1038/s41467-020-17236-y>
- Bohté SM, Kok JN, Poutré HL (2000) Spikeprop: backpropagation for networks of spiking neurons. In: *The European Symposium on Artificial Neural Networks*, URL <https://api.semanticscholar.org/CorpusID:14069916>
- Cohen G, Afshar S, Tapson J, et al (2017) Emnist: Extending mnist to handwritten letters. In: *2017 International Joint Conference on Neural Networks (IJCNN)*, pp 2921–2926
- Comsa IM, Potempa K, Versari L, et al (2020) Temporal coding in spiking neural networks with alpha synaptic function: Learning with backpropagation
- Crafton B, Parihar A, Gebhardt E, et al (2019) Direct feedback alignment with sparse connections for local learning. *Frontiers in neuroscience* 13:525
- Cramer B, Stradmann Y, Schemmel J, et al (2022) The heidelberg spiking data sets for the systematic evaluation of spiking neural networks. *IEEE Transactions on Neural Networks and Learning Systems* 33(7):2744–2757
- Davies M, Srinivasa N, Lin TH, et al (2018) Loihi: A neuromorphic manycore processor with on-chip learning. *IEEE Micro* 38(1):82–99. <https://doi.org/10.1109/MM.2018.112130359>

- Frady EP, Sanborn S, Shrestha SB, et al (2022) Efficient neuromorphic signal processing with resonator neurons. *Journal of Signal Processing Syst* 94(10):917-927. <https://doi.org/10.1007/s11265-022-01772-5>
- Furber SB, Galluppi F, Temple S, et al (2014) The spinnaker project. *Proceedings of the IEEE* 102(5):652-665. <https://doi.org/10.1109/JPROC.2014.2304638>
- Gerstner W, Kistler WM (2002) *Spiking Neuron Models: Single Neurons, Populations, Plasticity*. Cambridge University Press, <https://doi.org/10.1017/CBO9780511815706>
- Göltz J, Kriener L, Baumbach A, et al (2021) Fast and energy-efficient neuromorphic deep learning with first-spike times 3(9):823-835
- Han D, Yoo HJ (2019) Direct feedback alignment based convolutional neural network training for low-power online learning processor. In: 2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW), pp 2445-2452
- Hong C, Wei X, Wang J, et al (2020) Training spiking neural networks for cognitive tasks: A versatile framework compatible with various temporal codes. *IEEE Transactions on Neural Networks and Learning Systems* 31(4):1285-1296. <https://doi.org/10.1109/tnnls.2019.2919662>, URL <https://doi.org/10.1109/tnnls.2019.2919662>
- Huo Z, Gu B, Yang Q, et al (2018) Decoupled parallel backpropagation with convergence guarantee. [1804.10574](https://arxiv.org/abs/1804.10574)
- Jin Y, Zhang W, Li P (2018) Hybrid macro/micro level backpropagation for training deep spiking neural networks. Curran Associates Inc., Red Hook, NY, USA
- Kingma DP, Ba J (2017) Adam: A method for stochastic optimization. [1412.6980](https://arxiv.org/abs/1412.6980)
- Launay J, Poli I, Boniface F, et al (2020) Direct feedback alignment scales to modern deep learning tasks and architectures. In: *Proceedings of the 34th International Conference on Neural Information Processing Systems*. Curran Associates Inc., Red Hook, NY, USA, NIPS'20
- LeCun Y, Cortes C, Burges C (2010) Mnist handwritten digit database. ATT Labs [Online] Available: <http://yannlecom/exdb/mnist>
- Lee J, Zhang R, Zhang W, et al (2020) Spike-train level direct feedback alignment: Sidestepping backpropagation for on-chip training of spiking neural nets. *Frontiers in Neuroscience* 14. <https://doi.org/10.3389/fnins.2020.00143>
- Mostafa H (2016) Supervised learning based on temporal coding in spiking neural networks. *IEEE Transactions on Neural Networks and Learning Systems* PP
- Neftci EO, Augustine C, Paul S, et al (2017a) Event-driven random back-propagation: Enabling neuromorphic deep learning machines. *Frontiers in Neuroscience* 11:324

- Neftci EO, Augustine C, Paul S, et al (2017b) Event-driven random back-propagation: Enabling neuromorphic deep learning machines. *Frontiers in Neuroscience* 11. <https://doi.org/10.3389/fnins.2017.00324>, URL <https://doi.org/10.3389/fnins.2017.00324>
- Nøkland A (2016) Direct feedback alignment provides learning in deep neural networks. In: Lee D, Sugiyama M, Luxburg U, et al (eds) *Advances in Neural Information Processing Systems*, vol 29. Curran Associates, Inc.
- Refinetti M, D'Ascoli S, Ohana R, et al (2021) Align, then memorise: the dynamics of learning with feedback alignment. In: *International Conference on Machine Learning*, pp 8925–8935
- Rostami A, Vogginger B, Yan Y, et al (2022) E-prop on SpiNNaker 2: Exploring online learning in spiking RNNs on neuromorphic hardware. *Frontiers in Neuroscience* 16. <https://doi.org/10.3389/fnins.2022.1018006>, URL <https://doi.org/10.3389/fnins.2022.1018006>
- Rumelhart DE, Hinton GE, Williams RJ (1986) Learning representations by back-propagating errors. *nature* 323(6088):533–536
- Shrestha A, Fang H, Wu Q, et al (2019) Approximating back-propagation for a biologically plausible local learning rule in spiking neural networks. In: *Proceedings of the International Conference on Neuromorphic Systems*. Association for Computing Machinery, New York, NY, USA, ICONS '19, <https://doi.org/10.1145/3354265.3354275>
- Shrestha A, Fang H, Rider DP, et al (2021) In-hardware learning of multilayer spiking neural networks on a neuromorphic processor. In: *2021 58th ACM/IEEE Design Automation Conference (DAC)*, pp 367–372, <https://doi.org/10.1109/DAC18074.2021.9586323>
- Shrestha SB, Orchard G (2018) Slayer: Spike layer error reassignment in time. In: *Proceedings of the 32nd International Conference on Neural Information Processing Systems*. Curran Associates Inc., Red Hook, NY, USA, NIPS'18, p 1419?1428
- Takase H, Fujita M, Kawanaka H, et al (2009a) Obstacle to training spikeprop networks ? cause of surges in training process ?. pp 3062–3066, <https://doi.org/10.1109/IJCNN.2009.5178756>
- Takase H, Fujita M, Kawanaka H, et al (2009b) Obstacle to training spikeprop networks ? cause of surges in training process ?. In: *2009 International Joint Conference on Neural Networks*, pp 3062–3066, <https://doi.org/10.1109/IJCNN.2009.5178756>
- Wu Y, Deng L, Li G, et al (2018) Spatio-temporal backpropagation for training high-performance spiking neural networks. *Frontiers in Neuroscience* 12. <https://doi.org/10.3389/fnins.2018.00331>
- Wunderlich TC, Pehle C (2021) Event-based backpropagation can compute exact gradients for spiking neural networks. *Scientific Reports* 11(1)

Xiao H, Rasul K, Vollgraf R (2017) Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms

Zheng H, Wu Y, Deng L, et al (2020) Going deeper with directly-trained larger spiking neural networks. In: AAAI Conference on Artificial Intelligence, URL <https://api.semanticscholar.org/CorpusID:226290189>