



HAL
open science

EqSpike: Spike-driven equilibrium propagation for neuromorphic implementations

Erwann Martin, Maxence Ernoult, Jérémie Laydevant, Shuai Li, Damien Querlioz, Teodora Petrisor, Julie Grollier

► **To cite this version:**

Erwann Martin, Maxence Ernoult, Jérémie Laydevant, Shuai Li, Damien Querlioz, et al.. EqSpike: Spike-driven equilibrium propagation for neuromorphic implementations. *iScience*, 2021, 24 (3), pp.102222. 10.1016/j.isci.2021.102222 . hal-03451693

HAL Id: hal-03451693

<https://hal.science/hal-03451693v1>

Submitted on 15 Mar 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

EqSpike: Spike-driven Equilibrium Propagation for Neuromorphic Implementations

Erwann Martin¹, Maxence Ernoult^{2,3}, Jérémie Laydevant², Shuai Li², Damien Querlioz³, Teodora Petrisor¹, Julie Grollier^{2*}

¹ - Thales Research and Technology, 91767 Palaiseau, France

² - Unité Mixte de Physique, CNRS, Thales, Université Paris-Saclay, 91767 Palaiseau, France

³ - Université Paris-Saclay, CNRS, Centre de Nanosciences et de Nanotechnologies, 91120 Palaiseau, France

* julie.grollier@cnsr-thales.fr

Summary

Finding spike-based learning algorithms that can be implemented within the local constraints of neuromorphic systems, while achieving high accuracy, remains a formidable challenge. Equilibrium Propagation is a promising alternative to backpropagation as it only involves local computations, but hardware-oriented studies have so far focused on rate-based networks. In this work, we develop a spiking neural network algorithm called EqSpike, compatible with neuromorphic systems, which learns by Equilibrium Propagation. Through simulations, we obtain a test recognition accuracy of 97.6% on MNIST, similar to rate-based Equilibrium Propagation, and comparing favourably to alternative learning techniques for spiking neural networks. We show that EqSpike implemented in silicon neuromorphic technology could reduce the energy consumption of inference and training respectively by three orders and two orders of magnitude compared to GPUs. Finally, we also show that during learning, EqSpike weight updates exhibit a form of Spike Timing Dependent Plasticity, highlighting a possible connection with biology.

Introduction

Spike-based neuromorphic systems have, in recent years, demonstrated outstanding energy efficiency on inference tasks (Merolla *et al.*, 2014). Implementing the training of deep neural networks in such systems remains, however, a considerable challenge, as backpropagation does not apply directly to spiking networks and requires spatially non-local computations that go against the principles of neuromorphic systems. A large number of neuromorphic systems use the unsupervised and biologically-inspired Spike Timing Dependent Plasticity (STDP) learning rule because its weight updates, based on the relative timing of pre- and post-synaptic spikes, are spatially local and can be achieved with compact circuits in several technologies (Bi and Poo, 2001; Masquelier and Thorpe, 2007; Bichler *et al.*, 2012; Linares-Barranco *et al.*, 2011; Jo *et al.*, 2010; Pedretti *et al.*, 2017; Serb *et al.*, 2016; Prezioso *et al.*, 2018; Thakur *et al.*, 2018; Feldmann *et al.*, 2019). Unfortunately, STDP weight updates generally do not minimize a global objective function for the network, and the accuracy of STDP-trained neural networks remains below state-of-the-art algorithms based on the error backpropagation (Falez *et al.*, 2019). Important research efforts therefore investigate how the error backpropagation algorithm can be mathematically modified to make it spatially local and appropriate for spiking neural networks (Neftci *et al.*, 2017; Sacramento *et al.*, 2018; Richards *et al.*, 2019; Neftci, Mostafa and Zenke, 2019; Kaiser, Mostafa and Neftci, 2020; Bellec *et al.*, 2020; Payeur

et al., 2020). The derived learning rules are composed of three factors. The first two take into account, as usual, the behaviour of pre and post neurons, and the third allows for the introduction of an additional error factor. This third factor leads to implementations on neuromorphic chips that are less compact, and possibly less energy efficient, than two-factor learning rules such as STDP (Payvand *et al.*, 2020).

In this work, we propose a different approach to training spiking neural networks with high accuracy while using a local, two-factors learning rule compatible with neuromorphic implementations and scalable to complex tasks. Instead of starting from a non-local algorithm such as backpropagation and modifying it to make it local, we start from a rate-based algorithm called Equilibrium Propagation (Scellier and Bengio, 2017) that is intrinsically local in space, and features key advantages for neuromorphic implementations (Kendall *et al.*, 2020; Zoppo, Marrone and Corinto, 2020). Equilibrium Propagation theoretically applies to any physical system whose dynamics derive from an energy function. By minimizing the energy of such a system on data patterns, it can be made to relax towards states of minimal error prediction with respect to targets (Scellier and Bengio, 2017). The weight updates of Equilibrium Propagation match those of Back-Propagation-Through-Time (BPTT) in recurrent neural networks with static inputs (Ernout *et al.*, 2019), and it reaches high accuracy on image benchmarks such as CIFAR-10 (Laborieux *et al.*, 2020). Equilibrium Propagation uses the same set of weights for the forward and backward pass, a feature that is not biologically plausible, but is interesting for neuromorphic computing as it decreases the number of synaptic devices to update, thus reducing the overall power consumption. Contrarily to backpropagation, Equilibrium Propagation uses the same computations in the forward and backward phases, which is another highly desirable feature for neuromorphic systems as it greatly simplifies the circuits. Equilibrium propagation is, however, originally a rate-based algorithm.

Here, we design a spiking, hardware-friendly version of Equilibrium Propagation, called EqSpike, compatible with current neuromorphic technologies achieving online learning (Schemmel *et al.*, 2010; Furber *et al.*, 2014; Qiao *et al.*, 2015; Davies *et al.*, 2018; Frenkel *et al.*, 2019; Ishii *et al.*, 2019; Park, Lee and Jeon, 2020). EqSpike is local in space and time: contrarily to backpropagation, neither error gradients nor activations need to be stored in external memories, and synapses can be directly updated through neural events. We simulate a fully connected network based on this architecture on the MNIST handwritten digits database. We obtain a test recognition accuracy of 97.6%, which compares favourably with spiking neural networks learning with backpropagation-derived methods, and on par with rate-based Equilibrium Propagation. We show that EqSpike can be implemented in silicon neuromorphic technology, and thus reduce the energy consumption of inference by up to three orders of magnitude and training by up to two orders of magnitude compared to graphics processing units (GPUs). Finally, we also show that during learning the weight updates of EqSpike exhibit a form of STDP, yielding insights to its link to biology.

EqSpike: a hardware-friendly spiking version of Equilibrium propagation

Equilibrium Propagation is an algorithm for training convergent recurrent neural networks. Input neurons are clamped to a static input and all the other neurons, bi-directionally connected through synapses, evolve dynamically in time to reduce the energy of the network (Scellier and Bengio, 2017). The algorithm functions in two phases: a free phase and a nudging phase. In the free phase, performing inference, the network is let to reach equilibrium (Fig. 1a). Once this is done, inputs are kept clamped, and output neurons are nudged towards the desired output (Fig. 1b). During this nudging phase, the prediction error at the output layer is converted into a “force” acting upon output neurons and propagating to the rest of the system through time until a second equilibrium is reached. For training, synaptic values are updated by probing the neuron states after (Scellier and

Bengio, 2017) or during (Ernault *et al.*, 2020) the nudging phase through a learning rule that has been shown theoretically and numerically to match the updates of Back-Propagation Through Time, the state-of-the-art algorithm for such recurrent neural networks (Ernault *et al.*, 2019). It has been shown recently that Equilibrium Propagation also reaches accuracy within 1% of BPTT with convolutional architectures on the CIFAR-10 dataset (Laborieux *et al.*, 2020).

The original version of Equilibrium Propagation uses a rate-based formulation where dynamical neurons evolve smoothly in time. For a network of leaky integrate and fire neurons described by an Hopfield-like energy function $E(u) = 1/2 \sum_i u_i^2 - 1/2 \sum_{i \neq j} \rho(u_i) \rho(u_j) - \sum_i b_i \rho(u_i)$, where u are the membrane potentials of neurons and ρ their activation function, the Equilibrium Propagation learning rule is: $\Delta W_{ij} \sim (\rho_i \rho_j)_n - (\rho_i \rho_j)_f$, where the product $\rho_i \rho_j$ is measured at equilibrium, at the end of the nudge phase and the free phase (Scellier and Bengio, 2017). This rule can be extended to the case when weights are continuously updated during the nudging phase (Ernault *et al.*, 2020):

$$\frac{dW_{ij}}{dt} \sim \dot{\rho}_j \rho_i + \dot{\rho}_i \rho_j, \text{ (Eq.1)}$$

where W_{ij} is the synaptic weight connecting neurons i and j , and ρ_i, ρ_j are the rates of the two neurons. The network dynamics thus directly compute the error derivative, encoded in the rate derivative of the post neuron $\dot{\rho}$ and multiplied by the activation function of the pre neuron. The reformulation of this rate-based learning rule to a spiking neural network is therefore the following: *each time neuron i spikes, the weight should be updated by a quantity proportional to the derivative of the rate of neuron j , $\dot{\rho}_j$ (first term in Eq. 1), and reciprocally.*

We propose here a simple strategy, compatible with current electronic hardware, to implement this learning rule. It is illustrated in Fig. 1c in the form of a circuit, including spike detection elements at the output of each neuron, as well as dedicated blocks that extract the rate derivative from the spike trains of each neuron in real-time, in order to update synapses accordingly.

We use leaky-integrate-and-fire (LIF) spiking neurons which output spike frequency as a function of input current approximates the hard sigmoid prescribed in the original formulation of Equilibrium Propagation (Scellier and Bengio, 2017). Their maximum frequency is $f_{max} = 1/T_{refract}$, where $T_{refract}$ is the refractory time of the LIF neuron (see Supplementary Information for details and all parameter values).

The novelty compared to standard SNNs is the scheme that we propose for extracting the rate acceleration $\dot{\rho}$ for each neuron, illustrated in Fig. 1e. A leaky-integrator with a leak factor γ_{LI} (without reset nor spikes), takes as input the spike train emitted by the neuron to which it is connected and outputs a slowly varying signal proportional to the rate of the neuron spike train: $V_{LI} \sim \rho / \gamma_{LI}$ (Navarro *et al.*, 2020). To take the derivative, we delay this signal by a duration τ , and subtract the actual value with the delayed value:

$V_{delay} = V_{LI}(t) - V_{LI}(t - \tau) \cong \tau \frac{\partial V_{LI}}{\partial t} \propto \tau / \gamma_{LI} \dot{\rho}$. We then apply a low-pass filter for smoothing the variations. The filter is simulated using an average over N_{filt} simulation steps: $\overline{x(t)} = \frac{1}{N_{filt}} \sum_{i=0}^{N_{filt}-1} x_i(t - i dt)$, where dt is the simulation time step.

The output of the filter, that approximates $\frac{\tau}{\gamma_{LI}} \bar{\dot{\rho}}$, is then multiplied by the coefficient η_r . The corresponding weight updates are $\Delta w_{ij} = \eta_r \frac{\tau}{\gamma_{LI}} \bar{\dot{\rho}}_i$, which corresponds to an effective learning rate $l_r = \eta_r \frac{\tau}{\gamma_{LI}} = 1.5 \cdot 10^{-3}$.

This approach is hardware-compatible as LIF neurons, leaky integrators, delays and low pass filters are circuit elements that can be efficiently implemented in CMOS technology (Mead and Ismail, 1989), and bidirectional synapses could be implemented with CMOS compatible emergent nano-devices such as memristors (Ishii *et al.*, 2019; Marković *et al.*, 2020; Wan *et al.*, 2020). The corresponding pseudo-code is given in Algorithm 1 (see Supplementary Information for details).

Algorithm 1: EqSpike learning procedure for one image

Inputs: input image, Model $(n_{inputs}, n_{hidden}, n_{out})$, Loss function, length Free phase T_{free} , length Nudging phase T_{nudge} , Parameters $\gamma_{LIF}, \gamma_{LI}, u_{th}, \beta, \eta_r, \tau, N_{filt}, W_{ij}$

```

for  $t < T_{free}$ : ■ free phase
  for each neuron  $j$ :
    Update membrane potential  $u_j(\gamma_{LIF}, I_j)$ 
    if  $u_j > u_{th}$ :
      Emit a spike ( $t_j$ )
      Update  $\rho_j(t_j, \gamma_{LI})$ 

for  $t \in [T_{free}, T_{free} + T_{nudge}]$ : ■ nudging phase
  for each output neuron  $o$ :
    Compute error gradient  $\nabla e_o$ 
    Nudge neuron:  $u_o \leftarrow u_o - \beta \cdot \nabla e_o$ 
  for each neuron  $k$ :
    Update  $u_k(\gamma_{LIF}, I_k)$ 
    if  $u_k > u_{th}$ :
      Emit a spike ( $t_k$ )
    Update  $\rho_k(t_k, \gamma_{LI})$ 
    Compute smoothed:  $\bar{\rho}_k((\rho_k(t_k), \rho_k(t_k - \tau)), \dots, N_{filt})$ 

for each synapse  $w_{ij}$ : ■ Update synapses
  if neuron  $j$  emits a spike:
     $w_{ij} \leftarrow w_{ij} + \eta_r \cdot \frac{\tau}{\gamma_{LI}} \bar{\rho}_i$ 
  if neuron  $i$  emits a spike:
     $w_{ij} \leftarrow w_{ij} + \eta_r \cdot \frac{\tau}{\gamma_{LI}} \bar{\rho}_j$ 

```

Return: Trained weights for input image: W_{ij} and go to next image/next epoch.

Full network simulations: recognition rate on handwritten digits database

We now evaluate the performance of EqSpike on the MNIST handwritten digits classification task, using a fully connected network with one hidden layer (see Supplementary Information for details). The obtained train (orange) and test (blue) accuracies are shown Fig. 2 as a function of the number of training epochs, with the deviation over six runs in shadow color. Table 1 compares the results to BPTT and the version of Equilibrium Propagation closest to our implementation, called Continual Eq-Prop (Ernout *et al.*, 2020), trained with a batch size of one.

The test accuracy of EqSpike matches closely the accuracy of stochastic gradient descent through BPTT on the same network architecture, given the error margin. With a hidden layer of 300 neurons, EqSpikes reaches a test accuracy of 97.59%. Fully-connected spiking neural networks trained on MNIST without conversion from a non-spiking neural network typically achieve recognition rates in the 96-98% range (Neftci *et al.*, 2017; O’Connor and Welling, 2016; Lee, Delbruck and Pfeiffer, 2016; Mostafa, 2018; Tavanaei and Maida, 2019). EqSpike, with its local, two-factors online learning rule therefore reaches accuracies on MNIST comparable to those of the latest models investigated for training spiking neural networks on hardware platforms. We have chosen the MNIST dataset as a benchmark because it is a standard dataset for the neuromorphic community interested in training spiking neural networks online, due to the long simulation times. As EqSpike achieves results on MNIST equivalent to the baseline given by Equilibrium Propagation it has the potential, like Equilibrium Propagation, to adapt to convolutional architectures and perform with good accuracy on more complex image benchmarks (Laborieux *et al.*, 2020), with the additional advantage of being compatible with current neuromorphic technologies.

Inference speed and energy

As EqSpike is derived from a rate-based algorithm, it is interesting for neuromorphic applications to quantify the number of spikes needed to achieve inference, and the time needed to reach high accuracy. Operation with fewer spikes is more desirable, as it reduces both execution time and energy consumption.

Fig. 3 shows the inference results as a function of the execution time multiplied by the maximum frequency of neurons ($t \times f_{max}$). The orange line is the mean accuracy result over the whole test dataset, obtained by computing the spike rate of output neurons (as done for Fig. 2 and Table 1), through averaging in a time window $T_{average}$ of 100 simulation time steps ($T_{average} = 50/f_{max}$) and considering that the neuron with highest frequency encodes the output. This method computes the rate accurately at the expense of having to wait the time $T_{average}$ and letting output neurons spike multiple times. Spiking neural networks also offer the possibility to accelerate the computation and reduce the energy consumption by determining the output class from the first output neuron to spike. The blue line in Fig.3 is the accuracy as a function of time, averaged over all images in the test dataset, obtained by considering that the first output neuron to spike encodes the output.

The red, vertical dotted line in Fig. 3 indicates the average time of the first spike at the output over all presented images, corresponding to $t \cong 3.5/f_{max}$. At $t \cong 10/f_{max}$, the accuracy of single-spike inference (blue curve) reaches $95.11\% \pm 0.78\%$, within 1.4% of the precise rate computation (orange curve). This result shows that even though the algorithm is originally rate-based, a single spike at the output suffices in most cases to determine the correct class with good precision, a feature which is highly attractive for energy-efficient inference on neuromorphic chips. It means that inference can be achieved in $100 \mu s$ for electronic neurons with a firing rate of 100 kHz, available in neuromorphic chips working in accelerated time compared to biology (Schemmel *et al.*, 2010), and $1 \mu s$ for electronic neurons with a firing rate of 10 MHz that can be produced, for example, with emerging nanotechnologies (Li *et al.*, 2015). The corresponding throughputs are respectively 10k and 1M images/s, on par with current spiking neural network implementations (Pfeiffer and Pfeil, 2018; Park, Lee and Jeon, 2020). As the network operations are fully parallel, these orders of magnitudes will be conserved for wider networks. Simulations of rate-based Equilibrium Propagation on deeper networks indicate that the convergence time increases by a factor of about eight for a network with four hidden layers compared to a network of one hidden layer as here (Laborieux *et al.*, 2020). It

should be noted that in the current implementation we present static inputs to the network, which means that input neurons need to integrate these signals before they emit the first spikes that will then propagate to the next layers. The speed of inference could be increased in the future by presenting inputs directly encoded in spikes, for example sourced from neuromorphic vision sensors (Pfeiffer and Pfeil, 2018).

An estimation of the energy consumption of a spiking neural network on a neuromorphic silicon chip can be performed by counting the number of synaptic operations involved. Synaptic operations (SynOps) are defined as the total number of spikes transiting through synapses of the network. Frenkel et al show that a SynOp on a neuromorphic chip requires as little as 10 pJ (Frenkel *et al.*, 2019). The total number of synaptic operations needed for inference depends on the targeted recognition precision and, therefore, on the duration of inference (Fig. 3). For EqSpike, the recognition rate saturates at $t \cong 10/f_{max}$. The corresponding measured number of SynOps is about 150,000 in average. This is much less than expected if all neurons spiked. This is also a bit less but comparable to the number of SynOps needed at inference for Event-Driven Random BackPropagation (Neftci *et al.*, 2017). Considering 10 pJ/SynOps, each EqSpike inference could potentially consume 1.5 μ J. This means that testing the 10,000 images of the whole MNIST dataset could be achieved with a neuromorphic chip while consuming only 15 mJ, in other words, three orders of magnitude less than with a GPU (Joseph and Nagarajan, no date).

In our current EqSpike implementation, the input layer is the one leading to most spikes and SynOps: with only 16% of illuminated pixels in average in MNIST, the input layer emits 87.5% of all spikes and 98.6% of SynOps occur between the input layer and the hidden layer. In this work, we did not focus on reducing the number of spikes with encoding, but a better encoding of the input may reduce considerably the energy consumption. Kheradpisheh et al have shown that with a temporal encoding, the total number of spikes in the network before the first output spike can be reduced to 200 with a hidden layer four times larger than our network (Kheradpisheh and Masquelier, 2020). In our case, 678 spikes in total are emitted in average before the first output spike. An adaptation of EqSpike to temporal encoding is not straightforward, but this number could potentially be decreased in the future by reducing the encoding frequency of the input.

Training speed and energy

Training with EqSpike requires performing the free phase and then the nudging phase, during which synaptic weights are updated. A way to speed-up the training, and reduce the total number of SynOps, is to perform the nudging phase only on poorly classified examples, and skip the updates when the accuracy is satisfactory. We apply this strategy inspired from (Park, Lee and Jeon, 2020) using the criteria that the nudging phase is performed only when the difference between the target rate and the actual rate ($\widehat{\rho}_k - \rho_k$) at the output is above 1%. Fig. 4a shows the number of presented examples per epoch as a function of epoch number. In the last 20 epochs only approximately 15% of the training dataset still require a nudging phase.

For MNIST, we thus perform the free phase on all the dataset (3×10^6 images for the 50 epochs), and the nudging phase on 489,000 images. Given the durations of each phase, we can estimate the training time to $T_{training} \cong 2.74 \times 10^8 / f_{max}$. For electronic neurons with a firing rate of 100 kHz (Schemmel *et al.*, 2010), this leads to $T_{training} \cong 45$ min, and for electronic neurons with a firing rate of 10 MHz (Li *et al.*, 2015) to $T_{training} \cong 30$ s. As our networks feature a fully parallel nature, these training

times would be the same for much wider networks, and increased by a factor of about eight only with four hidden layers (Laborieux *et al.*, 2020).

As EqSpike is derived from a rate-based approach, it is interesting to compare the actual spiking rates of neurons in the network during training to their maximum frequency f_{max} . For neuromorphic applications, low overall rates are indeed desirable. Fig 4b shows the average number of spikes emitted by each neuron for an image presentation in the training dataset, as a function of the epoch. We found that for the training conditions of Fig. 2, there are in average 36 spikes/neuron/image. This means that neurons in the network spike in average with a frequency of the order of 20% of f_{max} , well below f_{max} , which is promising for neuromorphic implementations. Again this number could be reduced in the future by optimizing the encoding of input at the first layer.

Fig. 4c shows the numbers of synaptic operations needed for training as a function of recognition rate. The total number of synaptic operations after 50 epochs is of 4.23×10^{12} , which is of the same order of magnitude as Event-Driven Random Back-Propagation (Neftci *et al.*, 2017) for similar accuracy, and below training MNIST with BP based on (14). With 10pJ per SynOps (Frenkel *et al.*, 2019) the training phase of EqSpike on a neuromorphic chip could consume as little as 42 J, again, two orders of magnitude less than with a GPU (Joseph and Nagarajan, no date).

Spike Timing Dependent Plasticity

We have shown that EqSpike transforms Equilibrium Propagation into an efficient algorithm for neuromorphic chips. We now highlight that it also brings Equilibrium Propagation closer to biological plausibility. Bengio *et al* have pointed out a connection between the Equilibrium Propagation learning rule of Eq. 1 and STDP (Bengio *et al.*, 2017). The STDP learning rule, illustrated in Fig. 5a, reinforces causality between the spikes of pre- and post-synaptic neurons in networks with unidirectional synapses. If the post-synaptic neuron spikes after the pre-, causality is observed, and the weight is increased. In the opposite scenario, the weight is decreased.

Let us consider a situation where the pre-synaptic neuron i spikes and the post-synaptic neuron j accelerates, as illustrated in Fig. 5b. According to the Eq-Prop learning rule, in a network with unidirectional synapses, $\frac{dW_{ij}}{dt} \propto \dot{\rho}_j \rho_i = \dot{\rho}_{post} \rho_{pre}$, a positive weight update should be applied. Due to the acceleration of the post-neuron, there are less post-neuron spikes before the pre-neuron spike than after. Therefore, $t_{post} - t_{pre}$ is positive in average, yielding a positive weight update through STDP.

We have investigated if STDP-like weight updates did emerge during learning in our simulations. For this purpose, we monitored weight variations in synapses that connect input neurons and the hidden layer neurons. These synapses are unidirectional as input neurons are clamped to the input (see Fig. 1): their frequency does not vary. We used 100 images during the first epoch for the MNIST dataset. The curve in Fig. 5c shows the average weight updates in the first layer as a function of the time difference between post-synaptic neuron spikes and the average time of pre-synaptic neuron spikes in a window of 200 time steps before the post-synaptic neuron spike. The obtained curve, centred on zero, indeed exhibits an STDP-like shape. It has been obtained by filtering out very low frequencies below 0.05. Quiet neurons in the free phase indeed induce large weight updates at the beginning of the nudging phase, due to the sudden acceleration from zero to non-zero frequency, inducing an additional noise in the curve. It should be noted that biological STDP curves being frequency-dependent, they are also often obtained by focusing on a given range of frequencies (Kirkwood, Rioult and Bear, 1996). Fig. 5c shows that the STDP amplitude and time-window vary with the

strength of nudging β . In Equilibrium Propagation, weight changes are driven by neuron accelerations. As nudging applies a force that accelerates neurons, for larger β , larger weight changes are obtained, leading to a larger amplitude of the STDP-like curve. As higher accelerations are reached during the nudging phase, noticeable weight modifications occur for larger differences between t_{post} and t_{pre} , corresponding to a wider time window for the STDP-like curve. These results confirm the possible connection between STDP and Equilibrium Propagation pointed out in (Bengio *et al.*, 2017), despite the fact that individual spike timings are lost through averaging. They show that STDP-like behaviour can be obtained during learning in unidirectional synapses without a direct implementation of the original, causality-based rule. They also ask the question whether multilayer networks could be trained with the local STDP learning rule by using the nudging procedure of Equilibrium Propagation, possibly in the extended version called “Vector-field Equilibrium Propagation” in which synapses are unidirectional (Scellier *et al.*, 2018).

Discussion

Other versions of Equilibrium Propagation have been specifically designed to train spiking neural networks. O’Connor *et al* have designed a spiking network which, by construction, stochastically approximates the dynamics of its rate-based counter-part until reaching the same steady state with a minimal spikes communication budget between neurons (O’Connor, Gavves and Welling, 2019). Their technique, successfully tested against MNIST, comes at the cost of combining predictive coding, sigma-delta modulation and adaptive step sizes at the neuron level. In comparison, our technique simply requires LIF neurons, a spike-count and a low-pass filter to implement the learning rule. Mesnard *et al* proposed a version of Equilibrium Propagation to train spiking networks that also makes use of low-pass filters to estimate firing rates. However, they only demonstrate their approach on a non-linear toy problem, and the implemented learning rule is not local in time (Mesnard, Gerstner and Brea, 2016). In contrast, we demonstrate the effectiveness of our fully event-based implementation on MNIST.

More generally, several spike-based approaches to backpropagation have been proposed. One method consists in smoothing the spikes as a function of time so that gradients can be back-propagated through time (Huh and Sejnowski, 2018; Bellec *et al.*, 2020). Another technique consists in gating the spikes propagating the error signals by surrogate derivatives, as done in SpikeGrad (Thiele, Bichler and Dupret, 2019). Finally, Event-driven Random Backpropagation uses firing rates in the backward pass, as we do in this paper, and achieves similar performance on MNIST (Neftci *et al.*, 2017). In comparison with these approaches however, our implementation of Equilibrium Propagation does not require to know the activation function equation in order to compute its derivatives. One other interesting approach is S4NN, which employs latency coding, where each neuron can spike at most once, and the output is encoded as the first neuron to spike (Kheradpisheh and Masquelier, 2020). In contrast with rate-based coding, latency coding has the potential to save important energy in neuromorphic implementations. However, there is no clear indication of whether S4NN could scale to harder visual tasks, while rate-based Equilibrium Propagation was shown to train deep ConvNets on CIFAR-10 (Laborieux *et al.*, 2020). One intrinsic limitation of EqSpike, however, may stem from the necessity for the system to reach a steady state before the gradient computation phase. This prevents, for now, the classification of time-varying inputs. Payeur *et al* recently proposed a spike-based approach where top-down error signals are encoded as spike bursts and are multiplexed with bottom-up feedforward signals so that the backward pass and the forward pass can occur simultaneously with minimal disruption (Payeur *et al.*, 2020). While they show that the rate-based counterpart of their algorithm works on CIFAR-10 and ImageNet, it comes at the cost of employing dendritic network topologies and specialized synapses. Indeed, as noted in

the introduction, most backpropagation-derived local learning rules involve a third factor for supervision (Payvand *et al.*, 2020). In this regard, we believe that our EqSpike implementation of Equilibrium Propagation, with only two factors, achieves an optimal trade-off between circuitry complexity and performance.

Finally, Zoppo et al (Zoppo, Marrone and Corinto, 2020) and Kendall et al (Kendall *et al.*, 2020) have proposed using Equilibrium Propagation for training neuromorphic hardware. However, their implementations remain either rate-based or current-based, and are therefore not directly compatible with spiking neuromorphic chips. EqSpike, on the other hand, could be trained directly on reconfigurable neuromorphic systems with online learning such as SpiNNaker (Furber *et al.*, 2014) and Loihi (Davies *et al.*, 2018).

Limitations of the study

A limitation of EqSpike is that, with a wrong choice of γ_{LI} or τ combined with low nudging times, fast or brief rate changes cannot be detected, and the corresponding weight modification cannot be applied. This does not seem to be a problem with our experimentations as demonstrated above but could lead to a lower accuracy on more complex problems.

The circuits to compute the rate derivative could be further miniaturized by extracting $\dot{\rho}$ directly from the membrane potential if neuron models with smoothly varying membrane potential are used (Gerstner *et al.*, 2014).

EqSpike could also be sped up by building on dedicated hardware in analog or digital CMOS (Schemmel *et al.*, 2010; Qiao *et al.*, 2015; Thakur *et al.*, 2018; Frenkel *et al.*, 2019; Park, Lee and Jeon, 2020). Emerging nanotechnologies such as memristive synapses and nanoscale spiking oscillators are compelling candidates to scale up neuromorphic hardware due to their small size, their speed and their low energy consumption (Marković *et al.*, 2020; Milo *et al.*, 2020; Sebastian *et al.*, 2020; Wang *et al.*, 2020; Xi *et al.*, 2020). These technologies are typically prone to imperfections such as the device-to-device variability, cycle-to-cycle variability or the non-linearity in the conductance-to-voltage response, which are known to considerably jeopardize learning in memristive neural networks (Ishii *et al.*, 2019; Zhang *et al.*, 2020). Our paper implicitly assumes that the underlying memory technology at use would be linear, deterministic and identical across different synapses. The learning rate l_r in EqSpike is indeed about 10 times smaller than typically used in standard gradient descent. This is due to the fact that the batch size is one, and also that a low l_r is needed to avoid too large modifications accumulated over all the nudging phase. For low precision device as synapse, a first possible solution is to update the synaptic devices less frequently (not at each spike), but with a higher value. Another strategy is to adapt to EqSpike the training schemes used for binary neural networks (Hubara *et al.*, 2016; Rastegari *et al.*, 2016; Hirtzlin *et al.*, 2019). Further study should be done to propose a fully end-to-end circuit to implement EqSpike and investigate its resilience to the memristive device imperfections mentioned.

Conclusion

In this work we present a new algorithm for spiking neural networks, EqSpike, compatible with neuromorphic systems, and achieving good performance on MNIST. We show that EqSpike implements the learning rule of Equilibrium Propagation locally and autonomously. The gradients are computed by the dynamics of the system and the weights are modified by a spike and the addition of

only one block to the neuron. This can lead to spiking neuromorphic systems that do not need an external circuit to compute the error gradients given by backpropagation and learn autonomously, simply by presenting inputs and nudging the outputs according to errors. Our method obtained results on MNIST close to backpropagation through time and Equilibrium Propagation, two state-of-the-art algorithms. Moreover, because EqSpike is based on Equilibrium Propagation, the performance on more complex task, like CIFAR-10, could be similar. The number of synaptic operations to obtain these results in MNIST show we can obtain theoretically two orders of magnitude less energy consumption than a GPU for training and the same magnitude of time with high frequency neurons. The inference, after training the network, can be accelerated by waiting for the first output spike rather than compute the highest rate, with only a small loss of accuracy.

Finally, we show that the weight updates of EqSpike share similarity with STDP during learning, raising the question of a possible biological plausibility of the algorithm. In average, the modification of weight is proportional to the spike timing. This could permit to implement synapses in neuromorphic hardware by emergent nano-devices with a STDP-like behavior, to obtain a lower power consumption and higher surface density.

Author contributions

JG, TP and DQ have devised and supervised the study, EM has performed the spiking simulations with the help of ME, JL and SL. ME and JL have performed the benchmark rate-based simulations. All authors have contributed to writing the article.

Acknowledgements

This publication has received funding from the European Union's Horizon 2020 research innovation programme under grant agreement 732642 (ULPEC project) and 876925 (ANDANTE project)". E. Martin received funding from the ANRT grant 2018/0884.

Declaration of interests

The authors declare no competing interests.

Resource availability

Lead contact

Further information and requests for resources and reagents should be directed to and will be fulfilled by the lead contact, Julie Grollier (julie.grollier@cnrs-thales.fr).

Materials availability

N.A.

Data and code availability

The data generated by this study is available on reasonable request.

Figure 1 - EqSpike: spike-driven Equilibrium Propagation a) Schematic of the free phase in Equilibrium Propagation. b) Schematic of the nudging phase in Equilibrium Propagation. c) Illustration of the weight update implementation in EqSpike. d) Spiking rate of the neuron as a function of the amplitude of the input signal. e) Schematic of the rate acceleration computation.

Figure 2 - Recognition accuracy during training (orange) and test (blue) as a function of the number of epochs for MNIST, averaged over six runs.

Figure 3 - Inference time: recognition accuracy on MNIST on the test dataset as a function of time multiplied by the maximum neuron frequency f_{max} . Orange line: recognition accuracy computed from the output neuron showing the highest rate. Blue line: recognition accuracy computed from the output neuron spiking first. Red, vertical dotted line: average time of first output spike.

Figure 4 - Training performance. a) Number of presented images in the nudging phase per epoch versus epoch number. b) Number of spikes/neuron/image occurring during the two phases (nudge+free), as a function of the epoch. c) SynOps: number of spikes during both phases (nudge+free) as a function of the recognition accuracy.

Figure 5 - STDP. a) Illustration of the STDP learning rule; reproduction with data from (Bi and Poo, 2001). b) Illustration of the link between Eq-Prop and STDP learning rules; illustration reproduced from (Bengio et al., 2017). c) STDP-like curve during EqSpike learning.

Table 1 - Comparison between BPTT, C-EP and EqSpike, with the same initialization procedure. Batch size = 1.

References

Bellec, G. et al. (2020) 'A solution to the learning dilemma for recurrent networks of spiking neurons', *Nature Communications*, 11(1), p. 3625.

Bengio, Y. et al. (2017) 'STDP-Compatible Approximation of Backpropagation in an Energy-Based Model', *Neural Computation*, 29(3), pp. 555–577.

Bi, G. and Poo, M. (2001) 'SYNAPTIC MODIFICATION BY CORRELATED ACTIVITY: Hebb's Postulate Revisited', *Annual Review of Neuroscience*, 24(1), pp. 139–166.

Bichler, O. et al. (2012) 'Extraction of temporally correlated features from dynamic vision sensors with spike-timing-dependent plasticity', *Neural Networks*, 32, pp. 339–348.

- Davies, M. *et al.* (2018) 'Loihi: A Neuromorphic Manycore Processor with On-Chip Learning', *IEEE Micro*, 38(1), pp. 82–99.
- Ernault, M. *et al.* (2019) 'Updates of Equilibrium Prop Match Gradients of Backprop Through Time in an RNN with Static Input', in Wallach, H. *et al.* (eds) *Advances in Neural Information Processing Systems 32*. Curran Associates, Inc., pp. 7081–7091.
- Ernault, M. *et al.* (2020) 'Equilibrium Propagation with Continual Weight Updates', *arXiv:2005.04168 [cs, stat]*.
- Falez, P. *et al.* (2019) 'Unsupervised visual feature learning with spike-timing-dependent plasticity: How far are we from traditional feature learning approaches?', *Pattern Recognition*, 93, pp. 418–429.
- Feldmann, J. *et al.* (2019) 'All-optical spiking neurosynaptic networks with self-learning capabilities', *Nature*, 569(7755), pp. 208–214.
- Frenkel, C. *et al.* (2019) 'A 0.086-mm² 12.7-pJ/SOP 64k-Synapse 256-Neuron Online-Learning Digital Spiking Neuromorphic Processor in 28-nm CMOS', *IEEE Transactions on Biomedical Circuits and Systems*, 13(1), pp. 145–158.
- Furber, S. B. *et al.* (2014) 'The SpiNNaker Project', *Proceedings of the IEEE*, 102(5), pp. 652–665.
- Gerstner, W. *et al.* (2014) *Neuronal Dynamics: From Single Neurons to Networks and Models of Cognition*. 1 edition. Cambridge, United Kingdom: Cambridge University Press.
- Hirtzlin, T. *et al.* (2019) 'Hybrid Analog-Digital Learning with Differential RRAM Synapses', in *2019 IEEE International Electron Devices Meeting (IEDM)*. *2019 IEEE International Electron Devices Meeting (IEDM)*, p. 22.6.1-22.6.4.
- Hubara, I. *et al.* (2016) 'Binarized Neural Networks', in Lee, D. D. *et al.* (eds) *Advances in Neural Information Processing Systems 29*. Curran Associates, Inc., pp. 4107–4115.
- Huh, D. and Sejnowski, T. J. (2018) 'Gradient Descent for Spiking Neural Networks', in Bengio, S. *et al.* (eds) *Advances in Neural Information Processing Systems 31*. Curran Associates, Inc., pp. 1433–1443.
- Ishii, M. *et al.* (2019) 'On-Chip Trainable 1.4M 6T2R PCM Synaptic Array with 1.6K Stochastic LIF Neurons for Spiking RBM', in *2019 IEEE International Electron Devices Meeting (IEDM)*. *2019 IEEE International Electron Devices Meeting (IEDM)*, p. 14.2.1-14.2.4.
- Jo, S. H. *et al.* (2010) 'Nanoscale Memristor Device as Synapse in Neuromorphic Systems', *Nano Letters*, 10(4), pp. 1297–1301.
- Joseph, V. and Nagarajan, C. 'MADONNA: A Framework for Energy Measurements and Assistance in designing Low Power Deep Neural Networks.', p. 7.
- Kaiser, J., Mostafa, H. and Neftci, E. (2020) 'Synaptic Plasticity Dynamics for Deep Continuous Local Learning (DECOLLE)', *Frontiers in Neuroscience*, 14.
- Kendall, J. *et al.* (2020) 'Training End-to-End Analog Neural Networks with Equilibrium Propagation', *arXiv:2006.01981 [cs]*.
- Kheradpisheh, S. R. and Masquelier, T. (2020) 'Temporal Backpropagation for Spiking Neural Networks with One Spike per Neuron', *International Journal of Neural Systems*, 30(06), p. 2050027.

- Kirkwood, A., Rioult, M. G. and Bear, M. F. (1996) 'Experience-dependent modification of synaptic plasticity in visual cortex', *Nature*, 381(6582), pp. 526–528.
- Laborieux, A. *et al.* (2020) 'Scaling Equilibrium Propagation to Deep ConvNets by Drastically Reducing its Gradient Estimator Bias', *arXiv:2006.03824 [cs]*.
- Lee, J. H., Delbruck, T. and Pfeiffer, M. (2016) 'Training Deep Spiking Neural Networks Using Backpropagation', *Frontiers in Neuroscience*, 10.
- Li, S. *et al.* (2015) 'High-endurance megahertz electrical self-oscillation in Ti/NbOx bilayer structures', *Applied Physics Letters*, 106(21), p. 212902.
- Linares-Barranco, B. *et al.* (2011) 'On Spike-Timing-Dependent-Plasticity, Memristive Devices, and Building a Self-Learning Visual Cortex', *Frontiers in Neuroscience*, 5.
- Marković, D. *et al.* (2020) 'Physics for neuromorphic computing', *Nature Reviews Physics*, pp. 1–12.
- Masquelier, T. and Thorpe, S. J. (2007) 'Unsupervised Learning of Visual Features through Spike Timing Dependent Plasticity', *PLOS Comput Biol*, 3(2), p. e31.
- Mead, C. and Ismail, M. (eds) (1989) *Analog VLSI Implementation of Neural Systems*. Springer US (The Springer International Series in Engineering and Computer Science).
- Merolla, P. A. *et al.* (2014) 'A million spiking-neuron integrated circuit with a scalable communication network and interface', *Science*, 345(6197), pp. 668–673.
- Mesnard, T., Gerstner, W. and Brea, J. (2016) 'Towards deep learning with spiking neurons in energy based models with contrastive Hebbian plasticity', *arXiv:1612.03214 [cs, q-bio]*.
- Milo, V. *et al.* (2020) 'Memristive and CMOS Devices for Neuromorphic Computing', *Materials*, 13(1), p. 166.
- Mostafa, H. (2018) 'Supervised Learning Based on Temporal Coding in Spiking Neural Networks', *IEEE Transactions on Neural Networks and Learning Systems*, 29(7), pp. 3227–3235.
- Navarro, M. A. *et al.* (2020) 'Sodium channels implement a molecular leaky integrator that detects action potentials and regulates neuronal firing', *eLife*. Edited by K. J. Swartz, A. Mike, and B. P. Bean, 9, p. e54940.
- Neftci, E. O. *et al.* (2017) 'Event-Driven Random Back-Propagation: Enabling Neuromorphic Deep Learning Machines', *Frontiers in Neuroscience*, 11.
- Neftci, E. O., Mostafa, H. and Zenke, F. (2019) 'Surrogate Gradient Learning in Spiking Neural Networks: Bringing the Power of Gradient-Based Optimization to Spiking Neural Networks', *IEEE Signal Processing Magazine*, 36(6), pp. 51–63.
- O'Connor, P., Gavves, E. and Welling, M. (2019) 'Training a Spiking Neural Network with Equilibrium Propagation', in *The 22nd International Conference on Artificial Intelligence and Statistics. The 22nd International Conference on Artificial Intelligence and Statistics*, PMLR, pp. 1516–1523.
- O'Connor, P. and Welling, M. (2016) 'Deep Spiking Networks', *arXiv:1602.08323 [cs]*.

- Park, J., Lee, J. and Jeon, D. (2020) 'A 65-nm Neuromorphic Image Classification Processor With Energy-Efficient Training Through Direct Spike-Only Feedback', *IEEE Journal of Solid-State Circuits*, 55(1), pp. 108–119.
- Payeur, A. *et al.* (2020) 'Burst-dependent synaptic plasticity can coordinate learning in hierarchical circuits', *bioRxiv*, p. 2020.03.30.015511.
- Payvand, M. *et al.* (2020) 'Error-triggered Three-Factor Learning Dynamics for Crossbar Arrays', in *2020 2nd IEEE International Conference on Artificial Intelligence Circuits and Systems (AICAS). 2020 2nd IEEE International Conference on Artificial Intelligence Circuits and Systems (AICAS)*, pp. 218–222.
- Pedretti, G. *et al.* (2017) 'Memristive neural network for on-line learning and tracking with brain-inspired spike timing dependent plasticity', *Scientific Reports*, 7(1), p. 5288.
- Pfeiffer, M. and Pfeil, T. (2018) 'Deep Learning With Spiking Neurons: Opportunities and Challenges', *Frontiers in Neuroscience*, 12.
- Prezioso, M. *et al.* (2018) 'Spike-timing-dependent plasticity learning of coincidence detection with passively integrated memristive circuits', *Nature Communications*, 9(1), p. 5311.
- Qiao, N. *et al.* (2015) 'A reconfigurable on-line learning spiking neuromorphic processor comprising 256 neurons and 128K synapses', *Frontiers in Neuroscience*, 9.
- Rastegari, M. *et al.* (2016) 'XNOR-Net: ImageNet Classification Using Binary Convolutional Neural Networks', in Leibe, B. *et al.* (eds) *Computer Vision – ECCV 2016*. Cham: Springer International Publishing (Lecture Notes in Computer Science), pp. 525–542.
- Richards, B. A. *et al.* (2019) 'A deep learning framework for neuroscience', *Nature Neuroscience*, 22(11), pp. 1761–1770.
- Sacramento, J. *et al.* (2018) 'Dendritic cortical microcircuits approximate the backpropagation algorithm', in Bengio, S. *et al.* (eds) *Advances in Neural Information Processing Systems 31*. Curran Associates, Inc., pp. 8721–8732.
- Scellier, B. *et al.* (2018) 'Generalization of Equilibrium Propagation to Vector Field Dynamics', *arXiv:1808.04873 [cs, stat]*.
- Scellier, B. and Bengio, Y. (2017) 'Equilibrium Propagation: Bridging the Gap between Energy-Based Models and Backpropagation', *Frontiers in Computational Neuroscience*, 11.
- Schemmel, J. *et al.* (2010) 'A wafer-scale neuromorphic hardware system for large-scale neural modeling', in *Proceedings of 2010 IEEE International Symposium on Circuits and Systems. Proceedings of 2010 IEEE International Symposium on Circuits and Systems*, pp. 1947–1950.
- Sebastian, A. *et al.* (2020) 'Memory devices and applications for in-memory computing', *Nature Nanotechnology*, pp. 1–16.
- Serb, A. *et al.* (2016) 'Unsupervised learning in probabilistic neural networks with multi-state metal-oxide memristive synapses', *Nature Communications*, 7, p. ncomms12611.
- Tavanaei, A. and Maida, A. (2019) 'BP-STDP: Approximating backpropagation using spike timing dependent plasticity', *Neurocomputing*, 330, pp. 39–47.

Thakur, C. S. *et al.* (2018) 'Large-Scale Neuromorphic Spiking Array Processors: A Quest to Mimic the Brain', *Frontiers in Neuroscience*, 12.

Thiele, J. C., Bichler, O. and Dupret, A. (2019) 'SpikeGrad: An ANN-equivalent Computation Model for Implementing Backpropagation with Spikes', *arXiv:1906.00851 [cs]*.



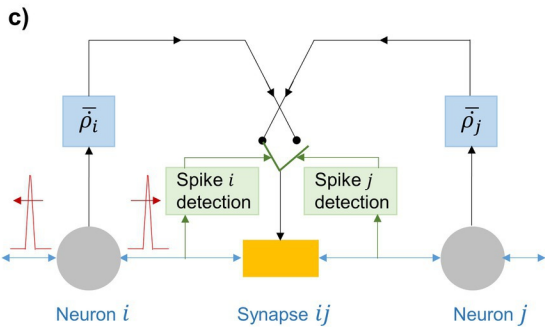


Wan, W. *et al.* (2020) '33.1 A 74 TMACS/W CMOS-RRAM Neurosynaptic Core with Dynamically Reconfigurable Dataflow and In-situ Transposable Weights for Probabilistic Graphical Models', in *2020 IEEE International Solid- State Circuits Conference - (ISSCC). 2020 IEEE International Solid- State Circuits Conference - (ISSCC)*, pp. 498–500.


Wang, Z. *et al.* (2020) 'Resistive switching materials for information processing', *Nature Reviews Materials*, 5(3), pp. 173–195.


Xi, Y. *et al.* (2020) 'In-Memory Learning With Analog Resistive Switching Memory: A Review and Perspective', *Proceedings of the IEEE*, pp. 1–29.




Zhang, W. *et al.* (2020) 'Neuro-inspired computing chips', *Nature Electronics*, 3(7), pp. 371–382.




Zoppo, G., Marrone, F. and Corinto, F. (2020) 'Equilibrium Propagation for Memristor-Based Recurrent Neural Networks', *Frontiers in Neuroscience*, 14.












	BPTT 784-100-10	Continual <u>Eq-Prop</u> 784-100-10	<u>EqSpike 100</u> 784-100-10	<u>EqSpike 300</u> 784-300-10
MNIST	Test: 97.11% \pm 0.23% Train: 99.06% \pm 0.15% (average on 6 runs)	Test: 96.97% \pm 0.12% Train: 99.8% \pm 0.04% (average on 6 runs)	Test: 96.87% \pm 0.18% Train: 98.59% \pm 0.03% (average on 6 runs)	Test: 97.59% \pm 0.1% Train: 98.91% \pm 0.03% (average on 6 runs)

EqSpike: Spike-driven Equilibrium Propagation for Neuromorphic Implementations

1- A spiking version of Equilibrium Propagation



3- Accuracy on MNIST reaches benchmark

Recognition rate: 97.6%

Fully connected architecture 784-300-10


Algorithm 1: EqSpike learning procedure for one image

```


Inputs: input image, Model ( $\gamma_{LIF}$ ,  $\beta$ ,  $\rho_{k=0}$ ,  $\rho_{k=1}$ ), Loss function, length free pl
Notifying phase  $T_{\text{notify}}$ , Parameters  $\gamma_{LIF}$ ,  $\beta$ ,  $\rho_{k=0}$ ,  $\rho_{k=1}$ ,  $\tau$ ,  $N_{\text{in}}$ ,  $N_{\text{h}}$ 
for  $t < T_{\text{train}}$ 
  for each neuron  $i$ 
    for each neuron  $j$ 
      Update membrane potential  $u_i(t_{\text{in}}-1)$ 
      if  $u_i > u_{\text{th}}$ 
        Emit a spike  $(t_i)$ 
        Update  $\rho_i(t_i, \rho_i)$ 
  for  $t \in [T_{\text{train}} - T_{\text{free}} + T_{\text{notify}}]$ 
    for each tested neuron  $i$ 
      Compute error gradient  $\nabla_{\rho_i}$ 
      Notify neuron:  $u_i = -\beta \cdot \nabla_{\rho_i}$ 
    for each neuron  $i$ 
      Update  $u_i(t_{\text{in}}-1)$ 
      if  $u_i > u_{\text{th}}$ 
        Emit a spike  $(t_i)$ 
        Update  $\rho_i(t_i, \rho_i)$ 
    Compute smoothed:  $\hat{\rho}_k((\rho_k(t_k), \rho_k(t_k - 1)), \dots, N_{\text{in}})$ 
  for each synapse  $W_{ij}$ 
    if neuron  $i$  emits a spike:
       $W_{ij} = W_{ij} + \eta \cdot \frac{\partial L}{\partial W_{ij}}$ 
    if neuron  $j$  emits a spike:
       $W_{ij} = W_{ij} + \eta \cdot \frac{\partial L}{\partial W_{ij}}$ 
  Return: Trained weights for input image:  $W_{ij}$  and go to next image/next epoch.
  
```



2- A hardware-compatible learning rule



4- STDP-like weight updates



5- High potential for hardware

Gain in energy consumption vs GPUs:

- Up to x 1000 for inference
- Up to x 100 for training