



An exact mapping from ReLU networks to spiking neural networks

Ana Stanojevic^{a,b,*}, Stanisław Woźniak^a, Guillaume Bellec^b, Giovanni Cherubini^{a,1}, Angeliki Pantazi^a, Wulfram Gerstner^b

^a IBM Research Europe – Zurich, Rüschlikon, Switzerland

^b École polytechnique fédérale de Lausanne, School of Life Sciences and School of Computer and Communication Sciences, Lausanne EPFL, Switzerland

ARTICLE INFO

Article history:

Received 24 January 2023
Received in revised form 31 August 2023
Accepted 4 September 2023
Available online 11 September 2023

Keywords:

Spiking neural network
ReLU network
Temporal coding
Single-spike network
Deep network conversion

ABSTRACT

Deep spiking neural networks (SNNs) offer the promise of low-power artificial intelligence. However, training deep SNNs from scratch or converting deep artificial neural networks to SNNs without loss of performance has been a challenge. Here we propose an exact mapping from a network with Rectified Linear Units (ReLUs) to an SNN that fires exactly one spike per neuron. For our constructive proof, we assume that an arbitrary multi-layer ReLU network with or without convolutional layers, batch normalization and max pooling layers was trained to high performance on some training set. Furthermore, we assume that we have access to a representative example of input data used during training and to the exact parameters (weights and biases) of the trained ReLU network. The mapping from deep ReLU networks to SNNs causes zero percent drop in accuracy on CIFAR10, CIFAR100 and the ImageNet-like data sets Places365 and PASS. More generally our work shows that an arbitrary deep ReLU network can be replaced by an energy-efficient single-spike neural network without any loss of performance.

© 2023 Published by Elsevier Ltd.

1. Introduction

Energy consumption of deep artificial neural networks (ANNs) with thousands of neurons poses a problem not only during training (Patterson et al., 2021; Strubell, Ganesh, & McCallum, 2020), but also during inference (Brown et al., 2020). Among other alternatives (Howard et al., 2017; Hubara, Courbariaux, Soudry, El-Yaniv, & Bengio, 2016; Tan & Le, 2019), hardware implementations of biologically inspired spiking neural networks (SNNs) (Attwell & Laughlin, 2001; Lennie, 2003) have been proposed as an energy-efficient solution (Burr et al., 2017; Davies et al., 2021; Diehl, Zarella, Cassidy, Pedroni, & Neftci, 2016; Gallego et al., 2020; Göltz et al., 2020, 2021; Sebastian et al., 2018), not only for large centralized applications, but also for computing in edge devices (Boroumand et al., 2021; Jiang, Chen, & Li, 2018; Wang et al., 2020). In SNNs neurons communicate by ultra-short pulses, called action potentials or spikes, that can be considered as point-like events in continuous time. In deep multi-layer SNNs, if a neuron in layer n fires a spike, this event causes a change in the voltage trajectory of neurons in layer $n + 1$. If, after some time, the trajectory of a neuron in layer $n + 1$ reaches a threshold value, then this neuron fires a spike.

While there is no general consensus on how to best decode spike trains in biology (Gerstner & Kistler, 2002; Pillow et al., 2008; Rieke, Warland, de Ruyter van Steveninck, & Bialek, 1996), multiple pieces of evidence indicate that immediately after an onset of a stimulus, populations of neurons in auditory, visual, or tactile sensory areas respond in such a way that the timing of the first spike of each neuron after stimulus onset contains a high amount of information about the stimulus features (Gollisch & Meister, 2008; Johansson & Birznieks, 2004; Kubke, Massoglia, & Carr, 2002). These and similar observations have triggered the idea that, immediately after stimulus onset, an initial wave of activity is triggered and travels across several brain areas in the sensory processing stream (Hung, Kreiman, Poggio, & DiCarlo, 2005; Optican & Richmond, 1987; Thorpe, Delorme, & Van Rullen, 2001; Thorpe, Fize, & Marlot, 1996; Yamins & DiCarlo, 2016). We take inspiration from these observations and assume in this paper that information is encoded in the exact spike times of each neuron and that spikes are transmitted in a wave-like manner across layers of a deep feedforward neural network.

Specifically, we use coding by time-to-first-spike (TTFS) (Gerstner & Kistler, 2002), a timing-based code originally proposed in neuroscience (Gerstner, 1998; Gerstner & Kistler, 2002; Gollisch & Meister, 2008; Johansson & Birznieks, 2004; Thorpe et al., 2001), which has recently attracted substantial attention in the context of neuromorphic implementations (Comsa et al., 2020; Gallego et al., 2020; Göltz et al., 2020, 2021; Kheradpisheh & Masquelier, 2020; Mostafa, 2018; Rueckauer & Liu, 2018; Stanojevic et al.,

* Correspondence to: Säumerstrasse 4, 8803 Rüschlikon, Switzerland.
E-mail address: ans@zurich.ibm.com (A. Stanojevic).

¹ Currently with Huawei Technologies – Zurich Research Center.

2022; Zhang et al., 2021). In our implementation of TTFS coding, each neuron fires exactly one spike. The stronger the input to a given neuron the earlier it fires. Coding schemes with at most one spike per neuron are intrinsically sparse in terms of the number of spikes. Since spike generation is a costly process from the energetic point of view (Sorbaro, Liu, Bortone, & Sheik, 2020), TTFS coding paves the way towards implementations with low energy consumption.

While a relation between ReLU networks and networks of non-leaky integrate-and-fire neurons with TTFS coding has been suggested before (Kheradpisheh & Masquelier, 2020; Mirsadeghi, Shalchian, Kheradpisheh, & Masquelier, 2021; Rueckauer & Liu, 2018; Zhang et al., 2021), there has been so far a major obstacle that prevented a successful exact mapping from an arbitrary deep ANN to a deep SNN with TTFS coding. In a standard ANN, time is discretized and layers are processed one after the other. The hard problem of an exact conversion of ReLU activations into spike times arises from the fact that in an SNN spikes are point-like events that arrive asynchronously in continuous time. Imagine a neuron in a ReLU network that receives several positive inputs that add up to a value of 0.8 and several negative inputs that add up to a value of -1.0 . Assuming a vanishing bias, the output of the unit is zero. However, if in the corresponding neuron of the SNN all the positive inputs have arrived *before* the negative inputs, the spiking neuron will have emitted a spike as soon as the firing threshold is reached (Rueckauer & Liu, 2018). Yet, it is impossible to “call back” the spike later on so as to cancel it. A potential solution to this problem is to consider that each layer starts its computation only once the calculation in the previous layer has finished. In a recent conversion scheme a time-dependent threshold was used to enforce the necessary waiting time (Rueckauer & Liu, 2018). However, due to imperfect conversion, there exists a small, but not negligible, loss in the final performance measure. Other conversion approaches for deep networks use custom activation functions (Bu et al., 2022) for ANN training, or numerically optimized multi-spike codes for a given activation function (Stockl & Maass, 2021). Moreover, most earlier approaches assumed that weights and biases of the ReLU network can be used as such in the corresponding SNN, but there is no fundamental reason why this should be the case.

What we would like for an exact mapping from ANN to SNN with TTFS coding is (i) a guarantee that no neuron fires too early so as to avoid the hard problem mentioned above; and (ii) a coding rule of how to translate the output $\bar{x}_i^{(n)}$ of neuron i in layer n of the ANN into a spike time of a corresponding neuron in the SNN. As an aside, we note that the hard problem of TTFS coding disappears if the trajectory of spiking neurons is always positive. The ideal mapping approach starts from a standard ReLU network with or without convolutional layers, batch normalization, and max pooling and maps it by a potentially nonlinear transformation of parameters to a corresponding SNN without any performance loss using a well-defined mapping rule from rates to spike timings.

In this paper we construct an explicit mapping that addresses the points above and guarantees the mathematical equivalence of an SNN with the corresponding ANN. We assume that there is a pretrained ReLU network and that we have access to its weights and biases as well as the input data on which the network was trained. Using TTFS coding, we propose a conversion which maps a deep ANN with ReLUs to an equivalent deep SNN with non-leaky integrate-and-fire units without any loss in performance. In contrast to other methods which require fine-tuning of the SNN or training an SNN from scratch (Bellec et al., 2020; Bohte, Kok, & La Poutre, 2002; Göltz et al., 2020; Kheradpisheh & Masquelier, 2020; Neftci, Mostafa, & Zenke, 2019; Park, Kim, Na, & Yoon, 2020; Tavanaei, Ghodrati, Kheradpisheh, Masquelier, & Maida,

2019; Woźniak, Pantazi, Bohnstingl, & Eleftheriou, 2020; Yan, Zhou, & Wong, 2021; Zenke & Ganguli, 2018; Zenke & Vogels, 2021), the goal of our work is to derive the final SNN from the pretrained deep ReLU network. Further numerical approximation or optimization steps (Bu et al., 2022; Rueckauer & Liu, 2018; Stockl & Maass, 2021) are not needed in our approach. The key to building an efficient TTFS conversion method is to derive an exact mathematical equivalence between an arbitrary deep ReLU network and the corresponding spiking network. Using standard pretrained models available online, we demonstrate conversion with 0% performance loss on CIFAR10 and CIFAR100 (Geifman, 2018) datasets as well as larger ImageNet-like (Russakovsky et al., 2015; Yang, Yau, Fei-Fei, Deng, & Russakovsky, 2022) datasets such as Places365 (Zhou, Lapedriza, Khosla, Oliva, & Torralba, 2017) and PASS (Asano, Rupprecht, Zisserman, & Vedaldi, 2021) without any training or fine-tuning.

2. Results

The subsection ‘Main theoretical result’ formulates the precise claim of a family of exact mappings from ANN to SNN. We then present the main ideas of the proof for one specific mapping scheme before we sketch alternative mapping schemes. Finally we test our mapping algorithm on benchmark datasets.

2.1. Main theoretical result

Definition (Deep ReLU Network). A Deep ReLU Network consists of $M \geq 1$ layers of hidden neurons with full or convolutional feedforward connectivity. Each neuron implements a rectified linear function $x_i^{(n)} = [a_i^{(n)}]_+$, where $[\]_+$ denotes rectification, and its activation variable $a_i^{(n)}$ is defined as:

$$a_i^{(n)} = \sum_j w_{ij}^{(n)} x_j^{(n-1)} + b_i^{(n)} \quad (1)$$

with weights $w_{ij}^{(n)}$ and a bias $b_i^{(n)}$. An upper index $n = 0$ refers to the input layer and the i th input is denoted with $x_i^{(0)}$. Optionally the network may also contain processing steps of max pooling and batch normalization.

Our aim is to map each neuron of the Deep ReLU network to an integrate-and-fire neuron in the SNN so that each neuron fires exactly once. A spike at time $t_j^{(n-1)}$ of neuron j in layer $n - 1$ generates a step current input with amplitude $J_{ij}^{(n)}$ into neuron i of layer n . The voltage trajectory of neuron i in layer n evolves according to

$$\frac{dV_i^{(n)}}{dt} = \alpha_i^{(n)} H(t - t_{\min}^{(n-1)}) + \sum_j J_{ij}^{(n)} H(t - t_j^{(n-1)}) + I_i^{(n)}(t) \quad (2)$$

where H denotes the Heaviside step function with $H(x) = 1$ for $x > 0$ or zero otherwise. The integration starts at time $t_{\min}^{(n-1)}$ with a slope $dV_i^{(n)}/dt = \alpha_i^{(n)} > 0$. The slope parameters $\alpha_i^{(n)}$, the weights $J_{ij}^{(n)}$, and the thresholds $\vartheta_i^{(n)}$ are parameters of the SNN. Neuron i in layer n may also receive an additional input $I_i^{(n)}(t)$. If the trajectory of $V_i^{(n)}$ crosses the threshold $\vartheta_i^{(n)}$ at time t from below, then $t = t_i^{(n)}$ is the firing time of neuron i in layer n . In our mapping, we use $I_i^{(n)}(t)$ to induce a short current pulse so as to trigger a spike at time $t_{\max}^{(n-1)}$ if neuron i has not fired before.

We claim that any Deep ReLU network can be mapped exactly to an SNN with integrate-and-fire neurons.

Theorem (Exact Mapping from ANN to SNN). Given the network parameters $\{w_{ij}^{(n)}, b_i^{(n)}\}$ of a Deep ReLU network that has been trained

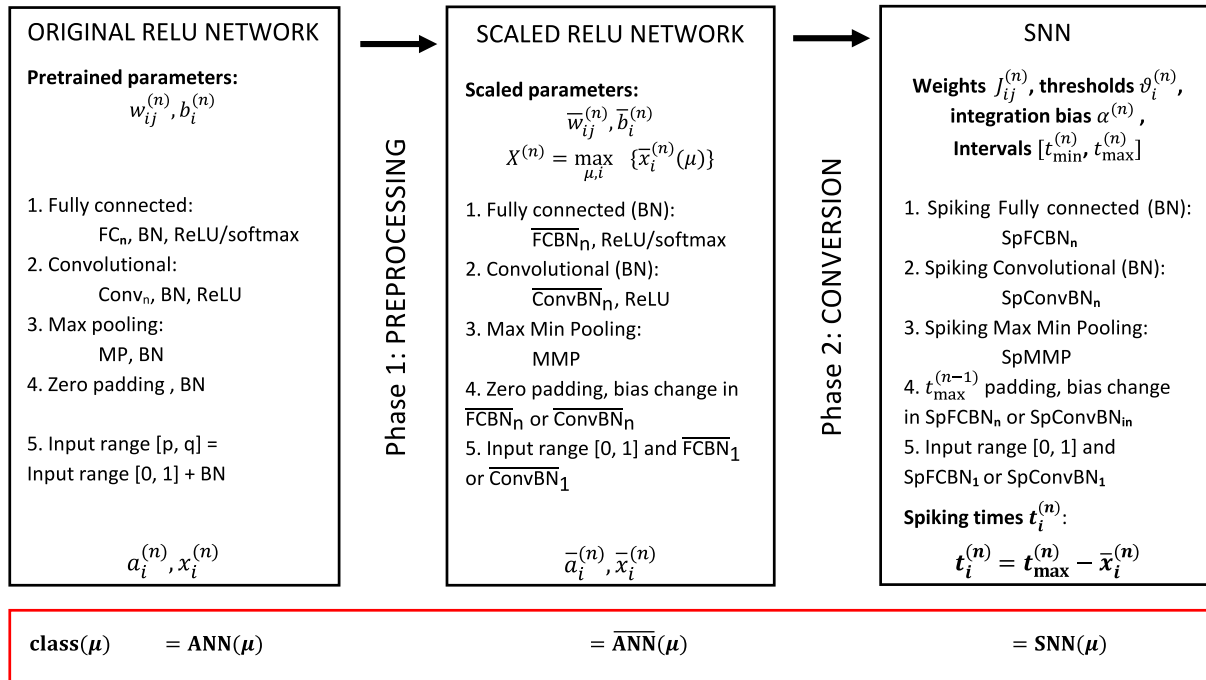


Fig. 1. Two phases of constructing a bidirectional mapping from ANN to SNN. In the first phase, Fully Connected (FC_n) or Convolutional ($Conv_n$) layers are identified; batch normalization (BN) is fused with the neighboring layers yielding layers \overline{FCBN}_n and \overline{ConvBN}_n , respectively; in case of zero padding, biases are adjusted at certain locations; max pooling steps are transformed to combined ‘Max Min Pooling’ (MMP) steps; and inputs are normalized to $[0, 1]$. Then the scaling symmetry of ReLU $x_i^{(n)} = [a_i^{(n)}]_+ = C[a_i^{(n)}/C]_+$ for an arbitrary constant $C > 0$ is applied to bring the weights of the ReLU network into a desired range and the maximum output $X^{(n)}$ for each layer is calculated. Overline indicates scaled ReLU network. In the second phase, the resulting parameters $\{\bar{w}_{ij}^{(n)}, \bar{b}_i^{(n)}, X^{(n)}\}$ are mapped to the parameters $\{J_{ij}^{(n)}, \vartheta_i^{(n)}, \alpha^{(n)}, t_{min}^{(n)}, t_{max}^{(n)}\}$ of the SNN. For an arbitrary input data point μ the three networks have the same values in the output layer (before applying softmax) and are therefore predicting the same class(μ).

to high performance on a training set and given access to a representative subset of the input data of the training set, there exists a family of bidirectional mappings from the Deep ReLU network to an SNN with TTFS coding without any loss in performance where each ReLU is replaced by an integrate-and-fire unit with dynamics as in Eq. (2) and parameters $\{J_{ij}^{(n)}, \vartheta_i^{(n)}, \alpha_i^{(n)}, t_{min}^{(n)}, t_{max}^{(n)}\}$.

Remarks. (i) The theorem mentions a family of mappings since the mapping is not unique, i.e., different combinations of parameters in the SNN give rise to an exact mapping. (ii) In the family of mappings that we consider each neuron emits at most a single spike. (iii) A consequence of the exact mapping is that both SNN and ANN have exactly the same performance on a sample-by-sample basis: if for a specific sample the prediction of the ANN is wrong then this is also the case for the SNN, and vice versa. (iv) One of the potential mappings is such that slope parameter $\alpha_i^{(n)}$ is identical for all neurons and all layers as stated in the following corollary.

Corollary (Mapping with Fixed α). An Exact mapping from ANN to SNN is possible with a slope parameter $\alpha_i^{(n)} = \alpha > 0$ that is identical for all neurons in all layers. Moreover, we may choose $t_{min}^{(n)} = t_{max}^{(n-1)}$.

2.2. Proof sketch of main theoretical result

Our proof is constructive, i.e., we propose an explicit mapping algorithm. The arguments work for arbitrary $\alpha_i^{(n)} > 0$. At the end of the argument we set $\alpha_i^{(n)} = \alpha$ to instantiate the conditions of the Corollary; see Methods for details. The algorithm has two phases, see Fig. 1.

Phase 1: Preprocessing. The original ReLU network undergoes preprocessing such that the network expects input in the $[0, 1]$

range and without any performance loss; furthermore batch normalization steps are removed by fusing them with the weights of neighboring layers; see Fig. 2, Algorithm 1 and Methods for details.

Importantly, and different to other studies in the field of network conversion, we use the known scaling symmetry of the ReLU activation function, i.e., $[a_i^{(n)}]_+ = C[a_i^{(n)}/C]_+$ for an arbitrary constant $C > 0$, to implement a nonlinear transformation from the original weight and bias parameters of the ReLU network to new parameters $\{\bar{w}_{ij}^{(n)}, \bar{b}_i^{(n)}\}$. After the transformation we can guarantee that the sum of the weights in each neuron is bounded in the range $(-B_{low}) \leq \sum_j \bar{w}_{ij}^{(n)} \leq 1 - \delta$ for hyperparameters $B_{low} > 0$ and $0 < \delta \ll 1$. The transformation proceeds layer-wise from the input to the output layer and does not change the network output. Neuron i in layer n of the scaled ReLU network has an activation variable $\bar{a}_i^{(n)}$ and output $\bar{x}_i^{(n)}$. The mapping is bidirectional so given the quantities $\bar{a}_i^{(n)}$ and $\bar{x}_i^{(n)}$ of the scaled network we can recover the variables $a_i^{(n)}$ and $x_i^{(n)}$ of the original network; see Methods Eqs. (20)–(23). Finally, since we have access to a representative sample of input data used during training, we extract $X^{(n)}$, the maximal activation of the rescaled ReLUs in layer n , across the input data and all neurons in layer n .

Phase 2: Conversion. To construct an exact conversion of the scaled ReLU network to the network of spiking neurons without any loss in performance we exploit six essential ideas (see Methods for details):

(i) *Choice of TTFS code.* We construct a mapping such that each neuron i in layer n of the SNN emits exactly one spike at $t_i^{(n)}$, where $t_{min}^{(n)} < t_i^{(n)} \leq t_{max}^{(n)}$ (see Fig. 3). Positive activation leading to a ReLU output $\bar{x}_i^{(n)} = \bar{a}_i^{(n)} > 0$ corresponds to an early firing time $t_i^{(n)} = t_{max}^{(n)} - \bar{a}_i^{(n)}$, or equivalently, $\bar{a}_i^{(n)} = t_{max}^{(n)} - t_i^{(n)}$. Thus,

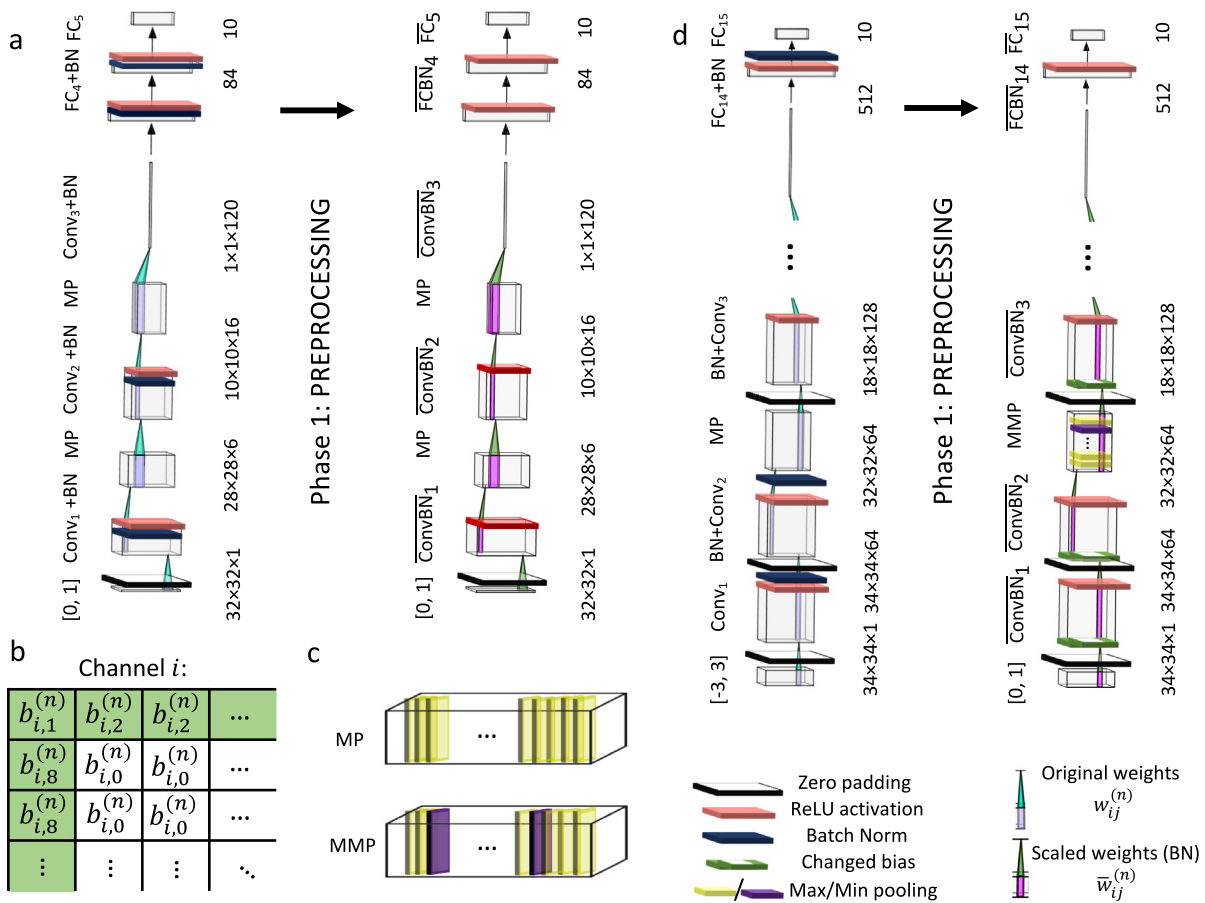


Fig. 2. Preprocessing for LeNet5 and VGG16 networks: a. LeNet5 original ReLU network with batch normalization (BN, black rectangular sheet) before the activation function (red sheet); during preprocessing batch normalization is fused with *previous* convolutional (Conv_{*n*}) and fully connected (FC_{*n*}) layers. b. When fusing a batch normalization layer with the next convolutional layer containing zero padding, some of the biases (in green) are changed; c. When fusing batch normalization with the next convolutional layer with max pooling (yellow) in between, specific channels might be changed to use min pooling function (violet). d. VGG16 original ReLU network used for CIFAR10 classification with batch normalization after the activation function; during preprocessing batch normalization is fused with *following* convolutional and fully connected layers.




Fig. 3. Mapping to multi-layer SNN with TTFS encoding and integrate-and-fire units. Left: Neurons in layer $n-1$ are connected to neuron i in layer n . Right: Spikes from neurons $1 \leq j \leq N^{(n-1)}$ arrive at times $t_j^{(n-1)}$ (green vertical arrows). The red trajectory shows the evolution of the potential $V_i^{(n)}$ as a function of time. Neuron i fires at time $t_i^{(n)}$ (red vertical arrow) when $V_i^{(n)}$ reaches the threshold $\vartheta_i^{(n)}$. The output value $\bar{x}_i^{(n)}$ of the corresponding ReLU corresponds to the time difference between $t_i^{(n)}$ and $t_{\max}^{(n)}$. Other neurons in layer n fire at other moments (blue, magenta). No neuron in layer n can fire later than $t_{\max}^{(n)}$. Our exact mapping procedure guarantees, that for all neurons the slope of the trajectory at the moment of spike firing is positive. Moreover, since in our theorem the threshold is arbitrarily high (thick blue vertical arrows) for $t < t_{\min}^{(n)}$, all firings in layer n occur for $t > t_{\min}^{(n)}$ so that the hard problem of late inhibitory input is solved.

spike times depend linearly on the output $\bar{x}_i^{(n)}$ of active ReLU neurons. Moreover if a neuron in layer n has not fired before time $t_{\max}^{(n)}$, it receives an additional external input pulse $I_i^{(n)}(t) = R \vartheta_i^{(n)} \delta(t - t_{\max}^{(n)})$ with $R \gg 1$ that triggers immediate firing at time $t_{\max}^{(n)}$. The parameters $\alpha_i^{(n)}$, $\vartheta_i^{(n)}$, the times $t_{\min}^{(n)} = t_{\max}^{(n-1)}$, as well as the weights $J_{ij}^{(n)}$ of the spiking network are determined during the conversion for all i, j, n as described in Methods (see Algorithm 2) and are kept fixed thereafter. With this coding scheme each neuron in layer $n - 1$ fires exactly once up to $t_{\max}^{(n-1)} = t_{\min}^{(n)}$. Therefore for $t > t_{\min}^{(n)}$ all input spikes to neurons in layer n have arrived.

(ii) *Slope of trajectory.* Since for $t > t_{\min}^{(n)}$ all input spikes to neurons in layer n have already arrived, the trajectory of neuron i in layer n has, a constant slope $dV/dt = \sum_j J_{ij}^{(n)} + \alpha_i^{(n)}$ [see. Eq. (2)] which is independent of the sequence of spike arrivals.

(iii) *Weight conversion.* Since for $t > t_{\min}^{(n)}$ the trajectories have constant slope, the mapping from activations in the ReLU to firing times in the SNN can be derived from the threshold-crossing condition $V_i^{(n)} = \vartheta_i^{(n)}$ for each neuron i in layer n . Evaluating this condition yields the nonlinear conversion of weights

$$J_{ij}^{(n)} = \frac{\alpha_i^{(n)}}{1 - \sum_{j'} \bar{w}_{ij'}^{(n)}} \bar{w}_{ij}^{(n)} \quad (3)$$

which is invertible. A similar invertible relation holds for the bias parameter (see Methods). Thus weights in the scaled ANN can be mapped to weights in the SNN without sign change. Summation over j on both sides of Eq. (3) shows that the slope has a value $\alpha_i^{(n)} + \sum_j J_{ij}^{(n)} = \alpha_i^{(n)} / (1 - \sum_j \bar{w}_{ij}^{(n)}) > 0$. Thus, once all input spikes have arrived, the slope of the trajectories is positive because of the weight rescaling $\sum_j \bar{w}_{ij}^{(n)} < 1$; This is the key motivation for the weight rescaling in Phase 1.

(iv) *Choice of $t_{\max}^{(n)}$.* Given our TTFS code, we know that a stronger activation leads to earlier spikes, yet we have to make sure that no neuron in layer n fires before the last spike of neurons in layer $n - 1$. The earliest possible spike in layer n occurs at time $t_{\max}^{(n)} - X^{(n)}$ where $X^{(n)}$ is the maximal activation of ReLU neurons in layer n identified during the preprocessing phase. We therefore set $t_{\max}^{(n)} = t_{\min}^{(n)} + (1 + \zeta)X^{(n)}$, where $\zeta > 0$. In practice (see below) a value of $\zeta = 0.5$ works well.

(v) *Choice of threshold.* By definition of our TTFS code, $t_{\max}^{(n)}$ is the time when a neuron in layer n that corresponds to a ReLU with activation $\bar{a}_i^{(n)} = 0$ reaches the threshold $\vartheta_i^{(n)}$; therefore this condition defines the value of the threshold. Because of different biases and different weights for different neurons, the thresholds $\vartheta_i^{(n)}$ are neuron-specific (see Methods). *This finishes the proof in the general case.*

(vi) *Free slope parameter.* Since the slope factor $\alpha_i^{(n)}$ is a free parameter, we can arbitrarily set $\alpha_i^{(n)} = 1$ for all neurons across all layers $1 \leq n \leq M$. This yields the Corollary. The condition of the Corollary is the specific case used in the simulations.

Remark. We may wonder how the above points solve the hard problem of TTFS coding. Our analysis above makes no statement about the trajectories of neurons in layer n for the time $t < t_{\min}^{(n)}$. Therefore we formally set the threshold for $t < t_{\min}^{(n)}$ to an arbitrarily high value to ensure that no spike occurs before $t_{\min}^{(n)}$. As mentioned under point (iii), our method guarantees a positive slope after $t_{\max}^{(n-1)} = t_{\min}^{(n)}$. Since during the allowed spiking interval $[t_{\min}^{(n)}, t_{\max}^{(n)}]$ the slope is fixed and positive, a spike never needs to be “called back”. Because of the preprocessing, we know that $\sum_j \bar{w}_{ij}^{(n)} > -B_{\text{low}}$. For example, a choice $B_{\text{low}} = 10$ and $\alpha_i^{(n)} = 1$ yields a slope larger than $1/11$. Furthermore, because of our choice of $t_{\max}^{(n)}$ under point (iv) we know that the interval $[t_{\min}^{(n)}, t_{\max}^{(n)}]$ is long enough to allow even the most

activated neuron to fire at the correct time. Finally, because of our choice of TTFS code under point (i) we are sure that all neurons in layer $n - 1$ have fired before or at $t_{\min}^{(n)}$. These choices together solve the hard problem of TTFS coding.

2.3. Examples of equivalent mappings

As stated in the main Theorem, the mapping from ANN to SNN is not unique; rather there is a family of equivalent mappings. Here we present several concrete implementation schemes.

2.3.1. Mapping with guaranteed positive slope

In the proof sketch above it was shown that the slope of all neurons is always positive once all input spikes have been received. However, we cannot exclude that before the time $t_{\min}^{(n)}$ the trajectory transiently has a negative slope; see Fig. 3. If this is desired for some application, we can use the free parameter $\alpha_i^{(n)} = \alpha^{(n)}$ to ensure that the slope of the trajectory is always non-negative, even before $t_{\min}^{(n)}$. To do so, we sum over all negative weights incoming to a given neuron and choose the slope parameter in layer n such that

$$\alpha^{(n)} + \min_i \sum_j J_{ij}^{(n)} H(-J_{ij}^{(n)}) > 0 \quad (4)$$

This ensures that the slope is positive not only if all inhibitory spikes arrive before the first excitatory spike, but also for all other possible timings of inhibitory input spikes. Therefore the hard problem of late inhibitory spikes can even be solved with a threshold that remains constant throughout the processing, i.e., even before $t_{\min}^{(n)}$. In practice we found that we could work with a constant threshold even if we did not implement the strict condition on the slope parameter formulated in Eq. (4) but worked instead with $\alpha^{(n)} = 1$. The strict condition in Eq. (4) can lead to very large slope parameters $\alpha^{(n)}$ and high firing thresholds $\vartheta_i^{(n)}$ which we might want to avoid in hardware implementations because of increased energy consumption.

2.3.2. Mapping with a dynamical threshold

In the proof sketch we assumed a constant threshold $\vartheta_i^{(n)}$ for all times $t > t_{\min}^{(n)}$. However, we can reinterpret the slope factor as a dynamical threshold. To see this, we integrate Eq. (2) and write the threshold condition that determines the firing time $t_i^{(n)} < t_{\max}^{(n)}$ in the form

$$\vartheta_i^{(n)} = V_i^{(n)}(t_i^{(n)}) = \alpha_i^{(n)} [t_i^{(n)} - t_{\min}^{(n-1)}] + \sum_j J_{ij}^{(n)} \epsilon(t_i^{(n)} - t_j^{(n-1)}) \quad (5)$$

where we have suppressed the external input and $\epsilon(s)$ is the voltage response to an input spike arriving at $s = 0$ (Gerstner & Kistler, 2002). Using standard textbook arguments, the term $\alpha_i^{(n)} [t_i^{(n)} - t_{\min}^{(n-1)}]$ can be shifted to the left-hand side which gives rise to a ‘dynamical threshold’ (Gerstner & Kistler, 2002) defined as $\vartheta_i^{(n)}(t) = \vartheta_i^{(n)} - \alpha_i^{(n)} [t - t_{\min}^{(n-1)}]$. Thus, the mapping in the corollary is identical to a mapping where the slope factor vanishes, but each neuron has a dynamical threshold that decreases linearly with time.

2.3.3. Mapping with identical weights in SNN and ANN

Previous studies have proposed approximative mappings under the condition $J_{ij}^{(n)} = \bar{w}_{ij}^{(n)}$ for all neurons in all layers. A quick glance at Eq. (3) tells us that a mapping with $J_{ij}^{(n)} = \bar{w}_{ij}^{(n)}$ becomes exact under the condition of a neuron-specific slope parameter

$$\alpha_i^{(n)} = 1 - \sum_j \bar{w}_{ij}^{(n)}. \quad (6)$$

Thus, in contrast to the mapping in the proof sketch of the Theorem, the slope parameter $\alpha_i^{(n)}$ is no longer a free parameter but *must* be chosen according to Eq. (6) if the aim is to have the same set of weights in ANN and SNN. Interestingly, under this condition, the trajectory of all neurons have the same slope of value one for $t > t_{\min}^{(n)}$.

2.3.4. Mapping with less than one spike per neuron

Even though our theory requires each neuron to spike exactly once, it is possible to have an alternative implementation where a given spiking neuron fires only when the corresponding ReLU is *active*. Instead of sending (costly) spikes of inactive neurons, it is sufficient to store the reference times $t_{\max}^{(n)}$ for all n . The trick is to set the slope of all trajectories of neurons i in layer n to $\alpha_i^{(n)} + \sum_j J_{ij}^{(n)}$ as soon as the maximum spike time $t_{\max}^{(n-1)}$ of neurons in layer $n - 1$ has been reached. This is mathematically equivalent to making all inactive neurons in layer $n - 1$ fire at time $t_{\max}^{(n-1)}$ but reduces the overall number of spikes in the network significantly. Therefore each neuron fires *at most* one spike. Since the ANN implements a nonlinear function from input to output, at least one neuron has to be inactive for at least one input data point, so that we know that on average there is strictly less than one spike per neuron. Since we are interested in a low-energy solutions, we report in the following the average number of active neurons across all inputs and all neurons for the given dataset. This number can be interpreted as ‘spikes per neuron per classification’. Note that this number depends on the specific regularization used during training of the ANN and can be further reduced by appropriate loss functions. Examples of this approach will be given below.

2.4. Performance on benchmark datasets

The above algorithm is a constructive proof that an exact mapping is possible. However, it is not clear how well it would perform in practice since there might be stability issues in the implementation or long processing delays that would reduce the attractivity of the mapping. In the following we test this algorithm on several image classification tasks with different standard datasets.

For each data set, we report the classification accuracy for the original ReLU network, for the SNN, as well as the percentage of agreement on a image-by-image basis between class prediction of original ReLU network and SNN network. Agreement is 100 percent, if for each image that is correctly (wrongly) classified by the ReLU, the image is also correctly (wrongly) classified by the SNN. Furthermore we report percentage of spikes per neuron under the implementation scheme mentioned at the end of the previous subsection.

2.4.1. MNIST, Fashion-MNIST and CIFAR10

In order to compare our results with existing conversion approaches (Rueckauer & Liu, 2018; Yan et al., 2021), we include MNIST and Fashion-MNIST (Deng, 2012; Xiao, Rasul, & Vollgraf, 2017) as well as CIFAR10 (Geifman, 2018) in our evaluation. We consider 16-layer VGG16, 5-layer LeNet5 and 2-layer fully connected networks, see Table 1. VGG16 contains max pooling, fully connected and convolutional layers together with zero padding and batch normalization applied after ReLU activation functions. For the MNIST dataset the SNN achieves the same 99.6% accuracy as the original ReLU network with 100% agreement, whereas the number of active neurons is around 51%. Similarly, for Fashion-MNIST there is a 100% agreement between SNN and ReLU predictions with the accuracy of 93.7% and around 45% of active neurons. In Rueckauer and Liu (2018) the authors perform a conversion of a 2-layer fully connected network as well as a

LeNet5 convolutional network such that the weights and biases of the SNN and ANN are identical. We reproduce the ANN results of those models and compare the performance of their SNN with the one obtained using our method. Our SNN surpasses the accuracy in Rueckauer and Liu (2018), and has 100% agreement between SNN and ANN with around 50% active neurons. Moreover, if we apply L1 activity regularization to the 2-layer fully connected network, MNIST images are classified with the accuracy of 98.52% with only 20 spikes in the hidden layer, i.e. 3.33% of active neurons. The original and the scaled LeNet5 network can be seen in Fig. 2a. For the preprocessing and mapping details refer to the Methods.

CIFAR10 contains color images of ten classes. The pretrained weights were obtained from an online repository (Geifman, 2018) where a convolutional network similar to the VGG16 architecture (see Table 1) was used. It comprises 15 layers in total since it uses only two fully connected layers instead of three (Fig. 2d). The network contains max pooling and convolutional layers together with zero padding and batch normalization applied after the ReLU activation functions. For CIFAR10 the SNN achieves the same 93.59% accuracy as the original ReLU network with 100% agreement between the two networks, whereas at 38% the number of active neurons is smaller than for the other datasets. If we apply L1 activity regularization, a small subset of 21.51% active hidden neurons is enough to perform classification with 93.16% accuracy. Further reduction of the number of spikes to 15.52% already decreases the accuracy to 90.54%.

In Fig. 4 we show an example of an SNN inference for classification of a cat image from CIFAR10. For input and hidden layers a raster plot of 10 neurons is shown and the spikes of neurons with higher activation of the corresponding neuron in the ANN are color-coded with darker shade. At the input layer the value of the data can be recovered from the spiking time of neuron i as $x_i^0 = 1 - t_i^{(0)}$ and in layer n the output of a neuron i of scaled ReLU network can be recovered from the spiking time of the corresponding neuron in the SNN as $\bar{x}_i^{(n)} = t_{\max}^{(n)} - t_i^{(n)}$. The duration of the interval $[t_{\min}^{(n)}, t_{\max}^{(n)}]$ varies considerably from one layer to the next. At the output layer $n = 15$ a potential with darker color indicates a larger value of the activation variable of the corresponding neuron in the original ReLU network. At time $t_{\max}^{(15)} = t_{\max}^{(14)} + 0.1$ when all the input from the layer $n = 14$ has arrived, the maximum potential corresponds to the neuron with maximal activation variable, i.e. both networks predict the same class.

2.4.2. Large-scale data sets

We avoided the ImageNet dataset because of privacy-concerns (Yang, Yau, Deng, & Russakovsky, 2022) and used instead Places365, PASS, and CIFAR100 for more realistic tests. The ‘Places365-Standard’ dataset contains high-resolution color images (Zhou et al., 2017) resembling those in the ImageNet dataset. The pretrained weights are obtained from an online repository (Zhou, 2018) that contains a standard VGG16 network without batch normalization which we map to a corresponding SNN; see Table 1. The SNN achieves the same 52.69% accuracy as the original ReLU network with 100% agreement between the two networks, whereas the number of active neurons is around 53%.

The PASS dataset consists of 1.4 million unlabeled images (Asano et al., 2021) and is used as a substitute for ImageNet (Asano et al., 2021) so as to avoid privacy-concerns. We use the same network as for the ‘Places365-Standard’ dataset, see Table 1. The weights are downloaded from the VGG16 model for ImageNet available in TensorFlow (tfp, 2022; Russakovsky et al., 2015). An inference on an image from PASS returns one of the 1000 ImageNet classes as output. When performing inference with the SNN

Table 1

Comparing performance of original ReLU networks and SNNs. Agreement metric shows percentage of inputs for which original ReLU network and SNN network predict the same class. Spikes column reveals the average percentage of active neurons across all hidden layers when the mapping with less than one spike per neuron is applied. Places365 and PASS are used as substitutes for ImageNet which breaches data privacy (Yang, Yau, Fei-Fei, et al., 2022). Accuracy calculation is not applicable for PASS dataset, as it is unlabeled. L1 denotes model with regularization. The highest accuracy and the lowest percentage of spikes for model/dataset pairs across different methods are highlighted in bold.

Model & dataset	Image size	Classes	Accuracy [%]		Agreement [%]	Spikes [%]
			ReLU	SNN		
Fully connected, MNIST (Rueckauer & Liu, 2018)	28 × 28 × 1	10	98.50	98.35	–	–
Fully connected, MNIST [ours]	28 × 28 × 1	10	98.52	98.52	100	50.28
Fully conn. L1, MNIST [ours]	28 × 28 × 1	10	98.52	98.52	100	3.33
LeNet5, MNIST (Rueckauer & Liu, 2018)	28 × 28 × 1	10	98.96	98.57	–	–
LeNet5, MNIST (Fig. 2a) [ours]	28 × 28 × 1	10	99.03	99.03	100	50.18
VGG16, MNIST [ours]	28 × 28 × 1	10	99.60	99.60	100	51.21
VGG16, Fashion-MNIST [ours]	28 × 28 × 1	10	93.70	93.70	100	45.34
VGG16, CIFAR10 (Yan et al., 2021)	32 × 32 × 3	10	92.55	92.48	–	–
VGG16, CIFAR10 (Fig. 2d) [ours]	32 × 32 × 3	10	93.59	93.59	100	38.38
VGG16 L1, CIFAR10 [ours]	32 × 32 × 3	10	93.16	93.16	100	21.51
Large-scale tests						
VGG16, CIFAR100 [ours]	32 × 32 × 3	100	70.48	70.48	100	38.21
VGG16, Places365 [ours]	224 × 224 × 3	365	52.69	52.69	100	53.72
VGG16, PASS [ours]	224 × 224 × 3	1000	N/A	N/A	100	53.24

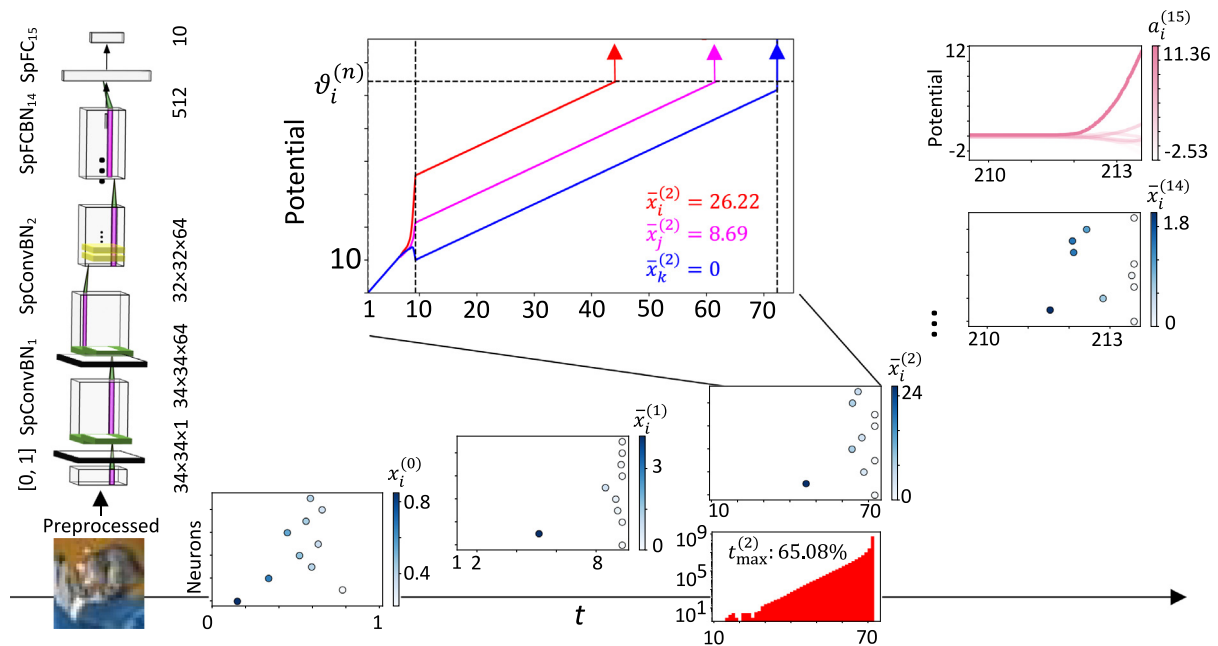


Fig. 4. Image Classification with spikes. A presentation of a cat image from the CIFAR dataset triggers spikes (filled and open circles) in consecutive layers of the SNN. For layers 0, 1, 2 and 14, one of the neurons is the one that fires the earliest spike for this image whereas the other nine are randomly selected. In a given layer, earlier spike times correspond to larger values (darker color) of the corresponding ReLU in the original network. At the output layer, the maximum potential corresponds to the largest activation variable in the ReLU network. Zoom inset: Voltage trajectories of three neurons from the same convolutional channel in layer 2. Spike times correspond to the moments of threshold crossing. Histogram inset: Histogram of spike counts per time bin of neurons in layer 2 averaged across all neurons and all images in the test set. Only 34.92% of neurons fire before $t_{\max}^{(2)}$.

we verify the agreement of the class prediction between the two networks. There is a 100% agreement between the original ReLU network and our SNN, with the fraction of active neurons around 53%. The results of this and the previous paragraph together show that the SNN achieves the same accuracy as the corresponding ANN on ImageNet-like datasets using spiking neurons that fire on average only for 53% of the inputs.

A similar statement is true for the CIFAR100 dataset. Using the same network architecture as for CIFAR10, and pretrained weights downloaded from an online repository (Geifman, 2018), we find on CIFAR100 a 100% agreement between SNN and ReLU predictions with the accuracy of 70.48% and around 38% of active neurons. Thus, on all tested large-scale datasets we find 100

percent agreement between the ANN and SNN indicating that the mapping is without any performance loss.

2.4.3. Sensitivity to noise and parameter changes

As outlined in the introduction, the hardest problem of the conversion is to prevent spike firing in layer n before all spikes from layer $n-1$ have arrived. In our mapping algorithm, a positive value $\zeta > 0$ should guarantee, for a large enough and representative subsample of input images from the training set, that during test the above problem is avoided. For all implementation results so far, the standard choice was $\zeta = 0.5$. In order to check sensitivity to the choice of ζ , we varied ζ across positive and negative values. Using the VGG16 model and the CIFAR10 dataset, we found that the performance degrades gracefully when

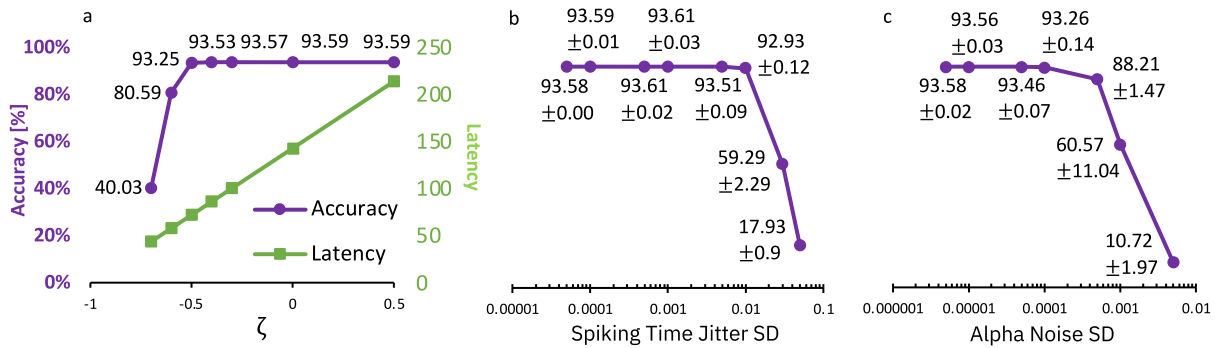


Fig. 5. Sensitivity tests a. Performance as a function of the parameter ζ . Note that the theorem requires $\zeta > 0$; our standard choice in all other figures is $\zeta = 0.5$. b. Performance of the network when there is a jitter in spike timing; note that spike timing difference between the earliest and latest neuron in a given layer is often less than 2 time units (see Fig. 4). c. Performance of the networks when neuron-specific variations are added to the slope parameter $\alpha_i^{(n)}$ during inference; note that the reference value is $\alpha^{(n)} = 1$ for all layers.

pushing ζ slightly into the negative regime, but breaks down for a value $\zeta < -0.5$, see Fig. 5a. Importantly, when switching from $\zeta = +0.5$ to $\zeta = -0.5$, the total processing time for image classification is reduced by a factor of three.

Noise in hardware implementations could potentially arise from a spike jitter caused for example by imprecisions in detecting the exact time of threshold crossing. We add a Gaussian noise of given standard deviation (SD) and perform 16 trials. No performance degradation is observed up to a standard deviation of 0.001, see Fig. 5b. With a jitter of about 1 percent, the accuracy drops from 93.59% to 92.93%, which depending on the application may or may not be considered as acceptable. We note that spike times of hundreds of neurons in a given hidden layer spread over an interval of one or a few time units so that even with a jitter of 0.01 the order of spike firing is considerably changed.

Imprecisions could also arise from heterogeneities in the hardware. A sensitive parameter is the reference slope $\alpha^{(n)}$. We modify the slope parameter in a neuron-specific way $\alpha_i^{(n)} + Y$ where Y is a zero-mean Gaussian random variable with a standard deviation that we control. This simulates a systematic neuron-specific hardware manufacturing imperfection. Even a standard deviation of 0.001 leads to a dramatic drop in performance, see Fig. 5c. This is expected since a small mismatch in slope leads to a relatively large shift in spike timing because changes are accumulated throughout the integration interval $[t_{\min}^{(n-1)}, t_{\max}^{(n)}]$. As mentioned in the discussion, using existing learning rules for spiking neurons in the hardware loop (Göltz et al., 2021) could be used to rapidly fine-tuning weights to compensate for hardware heterogeneities.

3. Discussion

In this paper we propose an exact mapping from a ReLU network to an SNN with time-to-first-spike coding. While a relation between ReLU networks and networks of non-leaky integrate-and-fire neurons has been suggested before (Kheradpisheh & Masquelier, 2020; Mirsadeghi et al., 2021; Rueckauer & Liu, 2018; Zhang et al., 2021), there have been four obstacles that in the past prevented a successful exact mapping from deep artificial neural networks to deep spiking neural networks:

(i) As mentioned in the introduction, a neuron in layer n that fires a spike before the last spike of neurons in the previous layer $n - 1$ has arrived could compromise an exact mapping, since not all inputs are taken correctly into account: in particular, a late inhibitory input could have led to substantially different spiking time if taken into account. Having access to a representative sample of inputs from the training data enables us to solve this problem by an appropriate choice of intervals $[t_{\min}^{(n)}, t_{\max}^{(n)}]$, with the

condition $t_{\max}^{(n-1)} = t_{\min}^{(n)}$. In other words, firing times of all neurons in layer n are guaranteed to fall into a desired interval, such that all spikes from layer $n - 1$ have arrived before the first neuron in layer n fires a spike.

(ii) In some implementations of an SNN, the slope of the potential of a neuron might be negative, zero, or only marginally positive once all input spikes have arrived. In the last case, the threshold could be eventually reached but spiking would be sensitive to noise. We have solved this problem by a positive slope parameter $\alpha_i^{(n)}$ for the trajectory of the integrate-and-fire neuron in combination with a suitable (non-unique) preprocessing of ReLU parameters that together guarantee that the slope of the trajectory is larger than some minimal value once all input spikes have arrived.

(iii) In the past it has been left open how to map the neuron of ReLU that is *inactive* for a given input vector to the corresponding spiking neuron. We have solved this problem by forcing the corresponding spiking neuron to fire a spike at the maximum spike time for that layer. We have also proposed an alternative implementation where inactive neurons do not fire spikes.

(iv) Existing spiking neural network approaches often use rate coding, custom activation functions or specific constraints during ANN training (Bellec et al., 2020; Gardner, Sporea, & Grüning, 2015; Huh & Sejnowski, 2018; Neftci et al., 2019; Rueckauer & Liu, 2018; Rueckauer, Lungu, Hu, Pfeiffer, & Liu, 2017; Woźniak et al., 2020; Yan et al., 2021; Zenke & Ganguli, 2018; Zenke & Vogels, 2021; Zhang, Zhou, Zhi, Du, & Chen, 2019). In contrast to prior work, our approach uses sparse TTFS coding, standard ANN elements and does not involve learning. The advantage of our approach in view of an application in neuromorphic edge devices is that a network consisting of standard fully connected and convolutional layers with ReLU activation function as well as max pooling and batch normalization can be pretrained using well-established optimization tools. After conversion, the SNN is guaranteed to have the exact same accuracy as the original ANN. The disadvantage is that hardware imperfections such as uncontrolled parameter variations are not taken into account during training.

TTFS coding for a conversion from ANN to SNN has been used before in an implementation (Rueckauer & Liu, 2018) that contains elements similar to our approach, but with a few important differences. First, we have a systematic way to define the end $t_{\max}^{(n)}$ of the allowed spiking interval. Second, we use a TTFS code with a *linear* relation between spike times and ReLU output whereas the relation is nonlinear in the earlier scheme (Rueckauer & Liu, 2018). Third, we identify for the case $\bar{w}_{ij}^{(n)} = J_{ij}^{(n)}$ an exact condition for the slope parameter and generalize to mappings where the weights are not simply copied from the ANN to the

SNN. The latter gives the freedom to choose the slope parameter so that the trajectory has *always* positive slope.

The success of our method paves the road to many future research direction including both theory and application:

(i) The discrete transition between spikes that are absent or present (depending on the input or on parameter variations) has plagued learning algorithms for spiking neural networks (Bohte et al., 2002; Göltz et al., 2020; Kheradpisheh & Masquelier, 2020; Neftci et al., 2019; Tavanaei et al., 2019; Zenke & Ganguli, 2018). Our theoretical contributions imply that spikes do not appear or disappear, but are rather shifted forward or backward within some finite interval. Earlier learning approaches have shown that those spikes that are triggered at moments when the slope of the potential is close to zero induce a high sensitivity of spike timing to small parameter changes. By introducing a positive slope parameter into an integrate-and-fire neuron in combination with a suitable (non-unique) preprocessing of ReLU parameters our mapping guarantees that the slope of the trajectory is at the moment of firing bounded within some favorable range, so that the problems of sensitivity or discrete transitions are avoided. Therefore, our mapping approach opens the path towards stable learning algorithms in single-spike deep SNNs, for example using the mapping suggested here as initialization of parameters (Stanojevic et al., 2023). Making a step towards increased biological plausibility the approach can also be extended to a leaky integrate-and-fire neuron model where each input spike causes a response described by a double-exponential filter as long as the coding interval $[t_{\min}^{(n)}, t_{\max}^{(n)}]$ is short compared to the two time constants of the exponential (Stanojevic et al., 2023).

(ii) Extension of the mapping to other architectures such as ResNet and to other types of neurons beyond non-leaky integrate-and-fire and ReLU would give the opportunity to have higher flexibility in terms of pretrained models. Moreover, it is of interest to further expand the theoretical framework such that it processes not just a single image but a stream of input data. This would present significant benefits for applications. At the moment, our method is limited to feedforward networks and an extension to Recurrent Neural Networks is left for future work.

(iii) To leverage our theoretical contribution for low-energy applications, a hardware implementation of this algorithm is desirable (Widmer et al., 2023). In view of future hardware implementations we are interested to further reduce the number of spikes and latency. With an improved implementation in combination with L1 regularization of the ANN, we have already reduced the fraction of spikes per neuron to well below 50%, see Table 1. The overall classification latency is defined as $t_{\max}^{(N)}$, which is the time instant when the readout at the output layer happens. The obtained values for $t_{\max}^{(N)}$ in VGG16 networks are around 3000 for MNIST and 2000 for Fashion-MNIST. The shallower networks have much smaller classification latency on the MNIST dataset of around 50 with LeNet5 and close to 15 for a 2-layer fully connected network. Using the VGG16 model, CIFAR10 is classified with a latency close to 200 and around 300 for CIFAR100. Larger datasets yield higher latency, with values around 50'000 for Places365 and close to 100'000 for PASS. We emphasize that the units are arbitrary. The classification latency can be reduced by a less conservative choice of meta-parameters of the mapping so as to reduce the dead time between spike arrival times in layers $n - 1$ and layer n . In particular a choice $\zeta = -0.5$ (instead of $\zeta = +0.5$) reduces the overall processing time by a factor of three without a dramatic loss in performance; see Fig. 5a. Further reduction of latency is achievable with $\zeta < -0.5$, as well as by decreasing the precision of spiking times. This is most effective when combined with retraining of weights and thresholds along the lines discussed above under point (i) (Stanojevic et al., 2023).

(iv) For hardware implementations the question of robustness to noise and heterogeneities is also important. We have started to explore the robustness of our algorithm to noise by adding a Gaussian noise of given standard deviation to the spike times of each layer, see Fig. 5b. Moreover, we have considered the case where the conversion was done with slope parameter $\alpha^{(n)} = 1$, while the hardware introduces fixed noise of given standard deviation for the slope parameter of each neuron, see Fig. 5c. It would be possible to fine-tune network weights with existing algorithm (Göltz et al., 2020) to compensate for hardware heterogeneities. Our current approach assumes asynchronous processing in continuous time and real-valued weights. In view of digital hardware implementations, it would also be of interest to study the effects of weight and time quantization. Future work on quantification of energy reduction will crucially depend on the concrete hardware implementation that is envisaged.

To summarize, this paper provides a constructive proof that deep ReLU networks and single-spike neural networks of integrate-and-fire neurons are mathematically equivalent. As a consequence, we reach functional deep spiking neural networks that have the same accuracy as ReLU networks and where spiking neurons fire at most one spike per neuron. Since spike transmission is an energy costly process in biology (Attwell & Laughlin, 2001) and neuromorphic hardware (Sorbaro et al., 2020), our mathematical results open a pathway to low-energy computing with deep neural networks.

4. Methods

4.1. Preprocessing

Before we perform the mapping from the ReLU network to the SNN, we perform a few preprocessing steps on the network with pretrained weights.

(i) If the network does not use batch normalization, this step is skipped. If batch normalization is implemented, it is fused into the neighboring fully connected and convolutional layers. The parameters of the batch normalization are $\hat{\mu}_i^{(n)}$ and $(\hat{\sigma}_i^{(n)})^2$ denoting the estimated mean and variance, $\gamma_i^{(n)}$ and $\beta_i^{(n)}$ which indicate scaling and shift factors learned during the optimization whereas ϵ is a small constant. In the following equations we use $\kappa_i^{(n)}$ to denote the scaling factor $\gamma_i^{(n)}/\sqrt{(\hat{\sigma}_i^{(n)})^2 + \epsilon}$.

When batch normalization is applied to the activation variable $a_i^{(n)}$ and before the activation function, it is fused with the processing of the previous layer (see Fig. 2a). The parameters are transformed as follows:

$$b_i^{(n)} \leftarrow \kappa_i^{(n)}(b_i^{(n)} - \hat{\mu}_i^{(n)}) + \beta_i^{(n)} \quad (7)$$

$$w_{ij}^{(n)} \leftarrow \kappa_i^{(n)} w_{ij}^{(n)}. \quad (8)$$

Note that in case of convolutional architecture each index i corresponds to a different channel.

When batch normalization is applied to the output of the activation function $x_i^{(n)}$, then it is fused with the processing of the subsequent layer (see Fig. 2d). The parameters are transformed as follows:

$$b_k^{(n+1)} \leftarrow b_k^{(n+1)} + \sum_i (\beta_i^{(n)} - \kappa_i^{(n)} \hat{\mu}_i^{(n)}) w_{ki}^{(n+1)}, \quad (9)$$

$$w_{ki}^{(n+1)} \leftarrow \kappa_i^{(n)} w_{ki}^{(n+1)}. \quad (10)$$

Note that the assignments of biases and weights need to be executed in this particular order. Moreover, in case of convolutional architecture, there are a few special cases that need to be considered.

Algorithm 1 Preprocessing

<p>Input: Model with parameters $w_{ij}^{(n)}$ and $b_i^{(n)}$</p> <p>Output: Model with parameters $\bar{w}_{ij}^{(n)}$ and $\bar{b}_i^{(n)}$</p> <p>1: $x_i^{(0)} \leftarrow \frac{x_i^{(0)} - p}{q - p}$</p> <p>2: for layer \in Model do</p> <p>3: if layer \in [Conv₁, FC₁] and $[p, q] \neq [0, 1]$ then</p> <p>4: Model \leftarrow fuse_BN_after_ReLU (Model, p, q)</p> <p>5: else if layer = BN and layer+1 = ReLU then</p> <p>6: Model \leftarrow fuse_BN_before_ReLU (Model)</p> <p>7: else if layer = BN and layer-1 = ReLU then</p> <p>8: Model \leftarrow fuse_BN_after_ReLU (Model)</p> <p>9: end if</p> <p>10: end for</p> <p>11: $c_i^{(0)} \leftarrow 0, n \leftarrow 1$</p> <p>12: for layer \in Model do</p> <p>13: if layer \in [Conv_n, FC_n] then</p> <p>14: Model, $c_i^{(n)} \leftarrow$ scale (Model, $c_i^{(n-1)}, B_{low}, \delta$)</p> <p>15: $X^{(n)} \leftarrow$ max_output ($\mu, \overline{\text{Model}}$)</p> <p>16: $n \leftarrow n + 1$</p> <p>17: end if</p> <p>18: end for</p>	<p>▷ Rescale inputs to [0, 1]</p> <p>▷ Iterate over all layers in the Model</p> <p>▷ Fuse imaginary batch normalization layer in case network was trained on $[p, q]$ range</p> <p>▷ Fuse batch normalization before activation with previous parametrized layer</p> <p>▷ Fuse batch normalization after activation with next parametrized layer and process padding or max pooling, Figs. 2b, 2c</p> <p>▷ Iterate over all layers in the Model</p> <p>▷ Scale layer with $c_i^{(n-1)}$, Eqs. (17), (19), and then with $c_i^{(n)}$, Eqs. (16), (18)</p> <p>▷ Calculate maximum output for layer n given training samples μ</p>
---	---

When batch normalization is applied to a zero-padded input into a convolutional layer, the bias change in Eq. (9) introduces an unnecessary offset at zero-padded locations. For these particular locations, we calculate the bias by taking into account only the set of inputs S_l which were not obtained through padding (see Fig. 2b). Eq. (9) is replaced with:

$$b_{k,l}^{(n+1)} \leftarrow b_{k,l}^{(n+1)} + \sum_{i \in S_l} (\beta_i^{(n)} - \kappa_i^{(n)} \hat{\mu}_i^{(n)}) w_{ki}^{(n+1)} \quad (11)$$

When max pooling is applied after batch normalization, the weights of the subsequent convolutional layer are changed as described in Eqs. (9) and (10). The batch normalization multiplies the output of each channel with factor $\kappa_i^{(n)}$, see Eq. (10). When this value is negative, the sign of the output is changed. During the inference time, the max pooling operation is transformed into a min pooling operation for the channels with switched sign (see Fig. 2c).

(ii) If network has input in range [0, 1], this step is skipped. Let us assume that the network has input in arbitrary range $[p, q]$. We would like for the network to operate for input in [0, 1] interval without changing its output. This scaling can be seen as an imaginary batch normalization layer between the input layer and the first layer.

The input data is transformed as $x_i^{(0)} \leftarrow \frac{x_i^{(0)} - p}{q - p}$ and the biases and weights of the first layer are set to:

$$b_k^{(1)} \leftarrow b_k^{(1)} + p \sum_i w_{ki}^{(1)}, \quad (12)$$

$$w_{ki}^{(1)} \leftarrow (q - p) w_{ki}^{(1)}. \quad (13)$$

When there is zero padding in the first convolutional layer, Eq. (12) is replaced with:

$$b_{k,l}^{(1)} \leftarrow b_{k,l}^{(1)} + p \sum_{i \in S_l} w_{ki}^{(1)} \quad (14)$$

(iii) In order to guarantee that the potential increases once all input spikes have arrived, we rescale the parameters of the ReLU network. We exploit the scaling symmetry of ReLU neurons $[a_i]_+ = C[a_i/C]_+$, for $C > 0$ and normalize weights so that the sum of input weights is smaller than $1 - \delta$, for some $0 < \delta < 1$. Similarly, we want to make sure that the sum of input weights does not fall below some lower bound $(-B_{low}) < 0$. To implement the scaling, we begin from the initial weights $\bar{w}_{ij}^{(n)} \leftarrow w_{ij}^{(n)}$ and biases $\bar{b}_i^{(n)} \leftarrow b_i^{(n)}$, start in layer $n = 1$ and proceed up to $n = M$ one layer at a time. For each neuron i , we calculate the sum over all the incoming weights

$$c_i^{(n)} = \sum_j \bar{w}_{ij}^{(n)} \quad (15)$$

If $c_i^{(n)} > (1 - \delta)$, we set for this specific neuron i its bias and incoming weights (for all j) to

$$\bar{b}_i^{(n)} \leftarrow \frac{(1 - \delta)}{c_i^{(n)}} \bar{b}_i^{(n)}; \quad \text{and} \quad \bar{w}_{ij}^{(n)} \leftarrow \frac{(1 - \delta)}{c_i^{(n)}} \bar{w}_{ij}^{(n)} \quad (16)$$

and the outgoing weights (for all k) to

$$\bar{w}_{ki}^{(n+1)} \leftarrow \frac{c_i^{(n)}}{1 - \delta} \bar{w}_{ki}^{(n+1)} \quad (17)$$

Similarly, if $c_i^{(n)} \leq -B_{low}$ we set the bias and the incoming weights (for all j) to

$$\bar{b}_i^{(n)} \leftarrow \frac{B_{low}}{|c_i^{(n)}|} \bar{b}_i^{(n)}; \quad \text{and} \quad \bar{w}_{ij}^{(n)} \leftarrow \frac{B_{low}}{|c_i^{(n)}|} \bar{w}_{ij}^{(n)} \quad (18)$$

and the outgoing weights (for all k) to

$$\bar{w}_{ki}^{(n+1)} \leftarrow \frac{|c_i^{(n)}|}{B_{\text{low}}} \bar{w}_{ki}^{(n+1)} \quad (19)$$

Note that signs are not changed by the scaling operation. Scaling ensures that for all hidden layers $(-B_{\text{low}}) \leq \sum_j \bar{w}_{ij}^{(n)} < 1$. We have larger weights in the final output layer (readout weights), but this does not cause any problems. The network where all the above preprocessing steps are applied is called a *scaled ReLU network*. Its parameters are denoted with a bar to distinguish them from the original, unscaled, network.

If $c_i^{(n)} > (1 - \delta)$ the activation variable $\bar{a}_i^{(n)}$ is given by

$$\bar{a}_i^{(n)} \leftarrow \frac{(1 - \delta)}{c_i^{(n)}} a_i^{(n)} \quad (20)$$

and if $c_i^{(n)} \leq -B_{\text{low}}$ by

$$\bar{a}_i^{(n)} \leftarrow \frac{B_{\text{low}}}{|c_i^{(n)}|} a_i^{(n)} \quad (21)$$

and $\bar{a}_i^{(n)} \leftarrow a_i^{(n)}$ otherwise. Similarly, if $c_i^{(n)} > (1 - \delta)$ the output $\bar{x}_i^{(n)}$ of ReLU is given by

$$\bar{x}_i^{(n)} \leftarrow \frac{(1 - \delta)}{c_i^{(n)}} x_i^{(n)} \quad (22)$$

and if $c_i^{(n)} \leq -B_{\text{low}}$ as

$$\bar{x}_i^{(n)} \leftarrow \frac{B_{\text{low}}}{|c_i^{(n)}|} x_i^{(n)} \quad (23)$$

and $\bar{x}_i^{(n)} \leftarrow x_i^{(n)}$ otherwise.

(iv) We apply all training data $1 \leq \mu \leq P$ at the input layer of the scaled ReLU network and observe the activation pattern for each neuron in the network. For each layer n we determine the maximal output of the activation function $\bar{x}_i^{(n)}(\mu)$ across all training data $1 \leq \mu \leq P$ and all neurons i in that layer:

$$X^{(n)} = \max_{i,\mu} \{\bar{x}_i^{(n)}(\mu)\} \quad (24)$$

If the number P is very large, we choose a statistically representative subset of data and perform the max-operation over these.

4.2. Conversion to SNN

The essential idea of the mapping from the ReLU neurons to the spiking neurons is that a positive activation leading to an output $\bar{x}_i^{(n)} = \bar{a}_i^{(n)} > 0$ is identified with an *early* firing time: $t_i^{(n)} = t_{\text{max}}^{(n)} - \bar{a}_i^{(n)}$, whereas vanishing output $\bar{x}_i^{(n)} = 0$ corresponds to firing at $t_i^{(n)} = t_{\text{max}}^{(n)}$.

The actual mapping is defined as follows (see Fig. 3).

(i) Input encoding. The input data lies in the interval $0 \leq x_i^{(0)} < 1$ and we set $t_{\text{min}}^{(0)} = 0$, $t_{\text{max}}^{(0)} = 1$ and $t_i^{(0)} = 1 - x_i^{(0)}$. With the parameters of the input layer fixed, we now proceed layer by layer from $n = 1$ to $n = M$

(ii) We set $t_{\text{min}}^{(n)} = t_{\text{max}}^{(n-1)}$

(iii) We set $t_{\text{max}}^{(n)} = t_{\text{min}}^{(n)} + B^{(n)}$ with $B^{(n)} = (1 + \zeta)X^{(n)}$ and $\zeta > 0$. The idea is that even the neuron with the strongest input must fire within the desired interval $[t_{\text{min}}^{(n)}, t_{\text{max}}^{(n)}]$, i.e., not too early. Under the assumption that the test data comes from the same statistical distribution as the training data, a small value $\zeta \ll 1$ should in practice provide a sufficient safety margin. Indeed, if the training data set is large enough to be statistically

representative, the probability that test data contains a point causing activation larger than $(1 + \zeta)X^{(n)}$ decreases rapidly with ζ . The range $[t_{\text{min}}^{(n)}, t_{\text{max}}^{(n)}]$ is therefore large enough to encode all the values from layer n of the rescaled ReLU network.

(iv) For a given $\alpha_i^{(n)} > 0$ we first choose a reference threshold $\tilde{v}_i^{(n)}$ in layer n such that an integrator without any spike input would fire at $t_{\text{min}}^{(n)}$. Hence for $t > t_{\text{min}}^{(n)}$ the reference threshold is

$$\tilde{v}_i^{(n)} = \alpha_i^{(n)} [t_{\text{max}}^{(n)} - t_{\text{min}}^{(n-1)}] \quad (25)$$

For the formal proof of the exact mapping, we set the reference threshold for $t \leq t_{\text{min}}^{(n)}$ to a sufficiently high value $\tilde{v}_i^{(n)} = \Theta \rightarrow \infty$ for all times $t \leq t_{\text{min}}^{(n)}$. This ensures that no neuron in layer n fires before $t_{\text{min}}^{(n)}$. The value from Eq. (25) is used only for $t > t_{\text{min}}^{(n)}$. However, for our practical algorithmic implementations we use the threshold given in Eq. (25) throughout for all t , because in all encountered data sets the probability of neurons in layer n firing before $t = t_{\text{min}}^{(n)}$ was negligible.

(v) The actual threshold also depends on the bias and weights of the neuron. To account for this, we set the actual threshold of neuron i in layer n to a value

$$v_i^{(n)} = \tilde{v}_i^{(n)} + D_i^{(n)} \quad (26)$$

With these parameter choices, an exact mapping from ReLU network to an SNN is possible with a value

$$D_i^{(n)} = [B^{(n)} \sum_j J_{ij}^{(n)}] - [\alpha_i^{(n)} + \sum_j J_{ij}^{(n)}] \bar{b}_i^{(n)}. \quad (27)$$

and weights

$$J_{ij}^{(n)} = \frac{\alpha_i^{(n)}}{1 - \sum_{j'} \bar{w}_{ij'}^{(n)}} \bar{w}_{ij}^{(n)} \quad (28)$$

where $\bar{w}_{ij}^{(n)}$ are the weights of the scaled ReLU network. Note that the denominator of Eq. (28) is always positive since $\sum_j \bar{w}_{ij}^{(n)} < 1$. Hence the mapping does not change the sign of the weights. The inverse weight transform from SNN to ReLU is

$$\bar{w}_{ij}^{(n)} = \frac{1}{\alpha_i^{(n)} + \sum_{j'} J_{ij'}^{(n)}} J_{ij}^{(n)} \quad (29)$$

This completes the conversion.

Note that we kept biases $\bar{b}_i^{(n)}$ as explicit parameters. However, following standard practice in the ANN literature, we could replace biases by an additional input neuron with connection weight equal to $\bar{b}_i^{(n)}$. The equations above as well as those for weight rescaling in Phase 1 are to be used analogously in that case.

Lemma. *With the conversion rules Eqs. (25) to (28) spike firing occurs at a value $\bar{x}_i^{(n)} = t_{\text{max}}^{(n)} - t_i^{(n)}$.*

Proof. Let us integrate the differential Eq. (2) of the integrate-and-fire units which yields for $t_{\text{min}}^{(n)} < t < t_{\text{max}}^{(n)}$ a voltage

$$V_i^{(n)}(t) = [t - t_{\text{min}}^{(n-1)}] \alpha_i^{(n)} + \sum_j J_{ij}^{(n)} [t - t_j^{(n-1)}], \quad (30)$$

where all neurons in layer $n - 1$ have firing times $t_j^{(n-1)} \leq t_{\text{max}}^{(n-1)} = t_{\text{min}}^{(n)}$. The firing time $t_i^{(n)}$ of neuron i in layer n is given by the threshold condition $V_i^{(n)}(t_i^{(n)}) = v_i^{(n)}$. We exploit that neurons in layer $n - 1$ that have not yet fired are forced to fire at $t_{\text{max}}^{(n-1)}$. We now insert the claims $t_i^{(n)} = t_{\text{max}}^{(n)} - \bar{x}_i^{(n)}$ and $t_j^{(n-1)} = t_{\text{max}}^{(n-1)} - \bar{x}_j^{(n-1)}$ into Eq. (30) and use Eqs. (25)–(27) as well as $t_{\text{min}}^{(n)} = t_{\text{max}}^{(n-1)}$ and

Algorithm 2 Conversion

Input: Model with parameters $\bar{w}_{ij}^{(n)}$ and $\bar{b}_i^{(n)}$, $X^{(n)}$

Output: SpModel with parameters $J_{ij}^{(n)}$, $\vartheta_i^{(n)}$, $\alpha^{(n)}$, $t_{\min}^{(n)}$ and $t_{\max}^{(n)}$

- 1: $t_{\min}^{(0)} \leftarrow 0$, $t_{\max}^{(0)} \leftarrow 1$, $n \leftarrow 1$, $\alpha^{(n)} \leftarrow 1$, $\forall n \leftarrow 1 \dots M$
- 2: **for** layer \in Model **do** ▷ Iterate layer-wise from the first to the output layer and calculate parameters and intervals
- 3: **if** layer \in [ConvBN_n, FCBN_n] **then**
- 4: **if** layer+1=ReLU **then**
- 5: $t_{\min}^{(n)} \leftarrow t_{\max}^{(n-1)}$
- 6: $t_{\max}^{(n)} \leftarrow t_{\max}^{(n-1)} + (1 + \zeta)X^{(n)}$
- 7: $J_{ij}^{(n)} \leftarrow \frac{\alpha^{(n)}\bar{w}_{ij}^{(n)}}{1 - \sum_j \bar{w}_{ij}^{(n)}}$
- 8: $\vartheta_i^{(n)} \leftarrow \alpha^{(n)}(t_{\max}^{(n)} - t_{\min}^{(n-1)}) + \sum_j J_{ij}^{(n)}(t_{\max}^{(n)} - t_{\min}^{(n)}) - (\alpha^{(n)} + \sum_j J_{ij}^{(n)})\bar{b}_i^{(n)}$
- 9: **else if** layer+1=softmax **then**
- 10: $\alpha_i^{(n)} \leftarrow \frac{\bar{b}_i^{(n)}}{(t_{\max}^{(n-1)} - t_{\min}^{(n-1)})}$
- 11: $J_{ij}^{(n)} \leftarrow \bar{w}_{ij}^{(n)}$
- 12: **end if**
- 13: $n \leftarrow n + 1$
- 14: **end if**
- 15: **end for**

$B^{(n)} = t_{\max}^{(n)} - t_{\min}^{(n)}$ to find

$$\bar{x}_i^{(n)} = \frac{1}{\alpha_i^{(n)} + \sum_j J_{ij}^{(n)}} \sum_j J_{ij}^{(n)} \bar{x}_j^{(n-1)} + b_i^{(n)} \quad (31)$$

Thus, the Eq. (29) for the weights follows from a comparison of this formula with the ReLU equation $\bar{x}_i^{(n)} = \sum_j \bar{w}_{ij} \bar{x}_j^{(n-1)}$; see Eq. (1) with $\bar{x}_i^{(n)} = [\bar{a}_i^{(n)}]_+$. The solution is unique since trajectories have positive slope so that the threshold is reached at most once.

4.3. Conversion of max pooling

If the ReLU network contains max pooling layers, the SNN contains layers performing max pooling and min pooling, outputting the earliest and latest spiking time respectively. This functionality can be implemented with integrate-and-fire neurons such that each neuron fires exactly one spike. To this end we introduce connections $K_{ij}^{(n-1)}$ within a given layer. A spike at time $t_j^{(n-1)}$ of a ReLU neuron j in layer $n-1$ generates a pulse current, modeled by a Dirac delta pulse of total charge $K_{ij}^{(n-1)}$, which is injected into neuron i of the max pooling or min pooling operation belonging to the layer $n-1$. The voltage of neuron i evolves according to

$$\frac{dV_{i(\text{MMP})}^{(n-1)}}{dt} = \sum_j K_{ij}^{(n-1)} \delta(t - t_j^{(n-1)}) \quad (32)$$

If $V_{i(\text{MMP})}^{(n-1)}$ crosses the threshold $\vartheta_{i(\text{MMP})}^{(n-1)}$ at time t then $t = t_{i(\text{MMP})}^{(n-1)}$ is the firing time of neuron i . For the layers which are preceded by a max pooling or min pooling operation the Eq. (2) is replaced with:

$$\frac{dV_i^{(n)}}{dt} = \alpha_i^{(n)} H(t - t_{\min}^{(n-1)}) + \sum_j J_{ij}^{(n)} H(t - t_{j(\text{MMP})}^{(n-1)}) + I_i^{(n)}(t) \quad (33)$$

In case of the max pooling operation, all weights $K_{ij}^{(n)}$ are set to slightly larger values than the threshold value $\vartheta_{i(\text{MMP})}^{(n)}$, such that the very first input spike triggers firing. In case of min pooling operation, parameters $K_{ij}^{(n)}$ are set to the value of $\vartheta_{i(\text{MMP})}^{(n)}/Q < K_{ij}^{(n)} < \vartheta_{i(\text{MMP})}^{(n)}/(Q-1)$ where Q is the total number of inputs. As a consequence, the very last input spike triggers the firing.

4.4. Mapping of the output layer

The output layer of the scaled ReLU network has a softmax activation function and parameters $\{\bar{w}_{ij}^{(M+1)}, \bar{b}_i^{(M+1)}\}$. In the SNN we implement the output layer with an integrator unit, i.e. the neurons just integrate the currents and do not spike. A spike arriving at the output layer at time $t_j^{(M)}$ from a neuron in layer M generates a step current input with amplitude $\bar{w}_{ij}^{(M+1)}$ into neuron i of layer $M+1$. The voltage of neuron i in layer $M+1$ evolves according to

$$\frac{dV_i^{(M+1)}}{dt} = \alpha_i^{(M+1)} H(t - t_{\min}^{(M)}) + \sum_j \bar{w}_{ij}^{(M+1)} H(t - t_j^{(M)}) \quad (34)$$

where H denotes the Heaviside step function. The non-leaky integration starts at time $t_{\min}^{(M)}$ and lasts until time $t_{\max}^{(M)}$ and $\alpha_i^{(M+1)}$ takes value:

$$\alpha_i^{(M+1)} = \frac{\bar{b}_i^{(M+1)}}{t_{\max}^{(M)} - t_{\min}^{(M)}} \quad (35)$$

The largest potential $V_i^{(M+1)}$ at time $t_{\max}^{(M)}$ determines the prediction.

4.5. Final remarks regarding the mapping

First, as mentioned in the results section, other mappings are also possible. For efficient coding with short latency, the aim is to choose parameters such that the resulting time intervals $[t_{\min}^{(n)}, t_{\max}^{(n)}]$ are not too large, however large enough to encode all values from the ReLU network with sufficient temporal resolution and such that the firing times of different layers do not overlap. Note that (in contrast to leaky integration with time-constant τ) a non-leaky integrator has no intrinsic time scale. Second, it would be possible to start the integration of all integrate-and-fire units across all layers n synchronously at time $t = 0$, if we increase at the same time the threshold in layer n by an amount $\alpha^{(n)} t_{\min}^{(n)}$.

Third, since $\alpha_i^{(n)} + \sum_j J_{ij}^{(n)} > 0$ and all neurons in layer $n-1$ have fired at or before $t_{\max}^{(n-1)}$, the voltage trajectories $V_i^{(n)}$ of all neurons i in layer n have for $t > t_{\max}^{(n-1)} = t_{\min}^{(n)}$ a positive slope; see.

Eq. (2). If, after preprocessing, $\sum_j \bar{w}_{ij}^{(n)} \leq 1 - \delta$, then the maximal slope of the trajectory at threshold is $\alpha^{(n)}/\delta$. Similarly, if after preprocessing $\sum_j \bar{w}_{ij}^{(n)} \geq -B_{low}$, then the minimal slope at the moment of firing is $\alpha^{(n)}/(1 + B_{low})$. A small slope of the potential close to the threshold should be avoided, since this increases the sensitivity to noise (in particular in view of combining with learning algorithms or unknown heterogeneities in the exact value of the slope). In practice, a value of $B_{low} = 10$ worked fine for our numerical simulations.

Fourth, we used a value of $\zeta = 0.5$. If the training set is large, and if we have access to all data in the training set, a positive but small $\zeta \rightarrow 0$ would be sufficient to guarantee that a neuron cannot fire ‘too early’. However, the test set could potentially include data where the total activation is slightly larger than the maximal activation in the training set. Since training set and test set arise, in principle, from the same statistical distribution, a parameter choice $\zeta = 0.5$ should provide a sufficient safety margin and this is confirmed in our simulations in the Results section.

4.6. Simulation details

All the experiments were performed with Python programming language and TensorFlow library. The simulations were executed on NVIDIA A100 GPUs.

In our algorithm we calculate the threshold crossing time $t_i^{(n)}$ exactly by solving Eq. (30) when $V_i^{(n)}(t_i^{(n)}) = \vartheta_i^{(n)}$. Therefore, our simulations work in continuous time with machine precision analogous to the real-valued output of the ReLU which is continuous with machine precision. Solving for exact threshold crossing times is a widely used technique in SNNs (Gewaltig & Diesmann, 2007; Mostafa, 2018) and conceptually different from a discrete-time implementation where threshold crossing is defined as the first time step where the voltage variable is above threshold.

We consider six datasets of different sizes and complexity:

(i) MNIST and Fashion MNIST datasets contain grayscale images of size 28×28 which are labeled into ten classes. For each of the two datasets there are 60000 training images and 10000 testing images. Data preprocessing step includes normalizing pixel values to the $[0, 1]$ range and in the case of a fully connected network the input is also reshaped. The pretrained parameters of the original ReLU networks are obtained by training with backpropagation using Adam optimizer (Kingma & Ba, 2015) with exponential learning rate schedule and standard cross-entropy loss. We apply dropout for regularization. In case of the VGG16 architecture the kernel was always of size 3 and the input of each convolutional operation is zero padded such that the shape at the output remains the same. Due to small input size, the first max pooling operation in the standard VGG16 architecture is omitted. The output of the convolutional part of VGG16 is of size 512 which is followed by two fully connected layers each containing 512 neurons and the output layer. The LeNet5 architecture has three convolutional, two max pooling and two fully connected layers with 84 and 10 neurons, see Fig. 2a. Finally, the 2-layer fully connected network has one hidden layer with 600 units. LeNet5 and VGG16-like networks also contain batch normalization before and after ReLU function, respectively.

In Fig. 2a we see the scaled LeNet5 network where the batch normalization are fused with previous convolutional and fully connected layers and the parameters of the network are scaled. For VGG16 network the batch normalization are fused with next convolutional and fully connected layers. Moreover, in this case the shift which appears due to zero padding is counter balanced with bias change at certain locations, see Fig. 2b, and every time batch normalization appears before max pooling, the channels whose sign is changed are replaced with min pooling, see Fig. 2c.

Since the model is trained on $[0, 1]$ range there is no need to fuse an imaginary batch normalization after the input. In order to obtain the scaled ReLU network the parameters of the network are scaled. Finding the maximum output $X^{(n)}$ of each layer on the subset of the training set finalizes the preprocessing step (see Fig. 1). In the following mapping phase the parameters of SNN are calculated.

(ii) CIFAR10 and CIFAR100 contain images of size $32 \times 32 \times 3$ (Krizhevsky & Hinton, 2009). For each of the two datasets there are 50000 training images and 10000 testing images. The data preprocessing step includes normalizing data with given fixed mean and standard deviation as given in Geifman (2018). The network was trained on the data rounded to $[-3, 3]$ range. In preparation for SNN mapping and inference, the input $x_i^{(0)}$ is further preprocessed as $x_i^{(0)} \leftarrow \frac{x_i^{(0)} + 3}{6}$. The kernel is always of size 3 and the input of each convolutional operation is zero padded such that the shape at the output stays the same. The output of the convolutional part of the VGG16 architecture has size 512 which is followed by two fully connected layers with 512 and 10 neurons.

During the preprocessing, the batch normalization is fused with the next convolutional and fully connected layers and bias is changed in locations where the input is coming from the zero padding. When necessary, the max pooling function is replaced with min pooling. Since the model is trained on $[-3, 3]$ range the imaginary batch normalization is fused with first convolutional layer and in locations where the input is generated by zero padding the bias is changed. In order to obtain the scaled ReLU network the parameters of the network are scaled. Finding the maximum output $X^{(n)}$ of each layer on the subset of training set finalizes the preprocessing step (see Fig. 1). In the following mapping phase the parameters of SNN are calculated.

(iii) The images in Places365-Standard dataset are labeled into 365 scene categories. There are 1.8 million training images, 36500 validation images and 328500 test images. Since the labels for the test set are not publicly available, we report the metrics on the validation set. Data preprocessing step includes centralizing data around a given fixed mean and reshaping it to the size of $224 \times 224 \times 3$ as described in Zhou (2018). The network is trained on the data which can be rounded to $[-200, 200]$ interval. In preparation for SNN mapping and inference the input $x_i^{(0)}$ is further preprocessed as $x_i^{(0)} \leftarrow \frac{x_i^{(0)} + 200}{400}$. Since the model is trained on $[-200, 200]$ range the imaginary batch normalization is fused with first convolutional layer and in locations where the input is generated by zero padding the bias is changed. In order to obtain the scaled ReLU network the parameters of the network are scaled. Finding the maximum output $X^{(n)}$ of each layer on the subset of training set finalizes the preprocessing step (see Fig. 1). In the following mapping phase the parameters of SNN are calculated.

(iv) We randomly sample 100000 testing and 5000 training images from PASS dataset. Most of the images in the dataset are colored and the few ones that are not are dropped during data preprocessing. The images are reshaped to $224 \times 224 \times 3$ and preprocessed with the same function as ImageNet for VGG16, which includes centering each color channel around zero mean. Since the model is trained on $[-200, 200]$ range, in preparation for SNN mapping and inference, the input $x_i^{(0)}$ is further preprocessed as $x_i^{(0)} \leftarrow \frac{x_i^{(0)} + 200}{400}$ situating the input on the $[0, 1]$ range. Moreover, the imaginary batch normalization is fused with the first convolutional layer and in locations where the input is generated by zero padding the bias is changed. In order to obtain the scaled ReLU network the parameters of the network are scaled. Finding the maximum output $X^{(n)}$ of each layer on the subset of training set finalizes the preprocessing step (see Fig. 1). In the following mapping phase the parameters of SNN are calculated.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The datasets utilized during the current study are available in public repositories.

Acknowledgments

The research of W.G. and G.B. was supported by a Sinergia grant (No CRSI15 198612) of the Swiss National Science Foundation.

References

- Applications - VGG16. (2022). URL https://www.tensorflow.org/api_docs/python/tf/keras/applications/vgg16/VGG16. (Accessed 21 November 2022).
- Asano, Y. M., Rupprecht, C., Zisserman, A., & Vedaldi, A. (2021). PASS: An ImageNet replacement for self-supervised pretraining without humans. *NeurIPS Track on Datasets and Benchmarks*.
- Attwell, D., & Laughlin, S. B. (2001). An energy budget for signaling in the grey matter of the brain. *Journal of Cerebral Blood Flow & Metabolism*, 21(10), 1133–1145.
- Bellec, G., Scherr, F., Subramoney, A., Hajek, E., Salaj, D., Legenstein, R., et al. (2020). A solution to the learning dilemma for recurrent networks of spiking neurons. *Nature Communications*, 11(1), 1–15.
- Bohte, S. M., Kok, J. N., & La Poutre, H. (2002). Error-backpropagation in temporally encoded networks of spiking neurons. *Neurocomputing*, 48(1–4), 17–37.
- Boroumand, A., Ghose, S., Akin, B., Narayanaswami, R., Oliveira, G. F., Ma, X., et al. (2021). Google neural network models for edge devices: Analyzing and mitigating machine learning inference bottlenecks. In *2021 30th International conference on parallel architectures and compilation techniques* (pp. 159–172). IEEE.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., et al. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33, 1877–1901.
- Bu, T., Fang, W., Ding, J., Dai, P., Yu, Z., & Huang, T. (2022). Optimal ann-snn conversion for high- accuracy and ultra-low-latency spiking neural networks. *ICLR*.
- Burr, G. W., Shelby, R. M., Sebastian, A., Kim, S., Kim, S., Sidler, S., et al. (2017). Neuromorphic computing using non-volatile memory. *Advances in Physics: X*, 2(1), 89–124.
- Comsa, I. M., Potempa, K., Versari, L., Fischbacher, T., Gesmundo, A., & Alakuijala, J. (2020). Temporal coding in spiking neural networks with alpha synaptic function. In *ICASSP 2020-2020 IEEE international conference on acoustics, speech and signal processing* (pp. 8529–8533). IEEE.
- Davies, M., Wild, A., Orchard, G., Sandamirskaya, Y., Guerra, G. A. F., Joshi, P., et al. (2021). Advancing neuromorphic computing with loihi: A survey of results and outlook. *Proceedings of the IEEE*, 109(5), 911–934.
- Deng, L. (2012). The mnist database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine*, 29(6), 141–142.
- Diehl, P. U., Zarella, G., Cassidy, A., Pedroni, B. U., & Neftci, E. (2016). Conversion of artificial recurrent neural networks to spiking neural networks for low-power neuromorphic hardware. In *2016 IEEE international conference on rebooting computing* (pp. 1–8). IEEE.
- Gallego, G., Delbrück, T., Orchard, G., Bartolozzi, C., Taba, B., Censi, A., et al. (2020). Event-based vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(1), 154–180.
- Gardner, B., Sporea, I., & Grüning, A. (2015). Learning spatiotemporally encoded pattern transformations in structured spiking neural networks. *Neural Computation*, 27(12), 2548–2586.
- Geifman, Y. (2018). Github. URL <https://github.com/geifmany/cifar-vgg>.
- Gerstner, W. (1998). Spiking neurons. In W. Maass, & C. M. Bishop (Eds.), *Pulsed neural networks* (pp. 3–53). MIT-Press, Ch. 1.
- Gerstner, W., & Kistler, W. K. (2002). *Spiking neuron models: Single neurons, populations, plasticity*. Cambridge UK: Cambridge University Press.
- Gewaltig, M., & Diesmann, M. (2007). NEST (neural simulation tool). *Scholarpedia*, 2(4), 1430.
- Gollisch, T., & Meister, M. (2008). Rapid neural coding in the retina with relative spike latencies. *Science*, 319(5866), 1108–1111.
- Göltz, J., Baumbach, A., Billaudelle, S., Kungl, A., Breitwieser, O., Meier, K., et al. (2020). Fast and deep neuromorphic learning with first-spike coding. In *Proceedings of the neuro-inspired computational elements workshop* (pp. 1–3).
- Göltz, J., Kriener, L., Baumbach, A., Billaudelle, S., Breitwieser, O., Cramer, B., et al. (2021). Fast and energy-efficient neuromorphic deep learning with first-spike times. *Nature Machine Intelligence*, 3, 823–835.
- Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., et al. (2017). Mobilenets: Efficient convolutional neural networks for mobile vision applications. arXiv preprint arXiv:1704.04861.
- Hubara, I., Courbariaux, M., Soudry, D., El-Yaniv, R., & Bengio, Y. (2016). Binarized neural networks. *Advances in Neural Information Processing Systems*, 29.
- Huh, D., & Sejnowski, T. J. (2018). Gradient descent for spiking neural networks. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, & R. Garnett (Eds.), *Advances in neural information processing systems: Vol. 31*, Curran Associates, Inc., URL <https://proceedings.neurips.cc/paper/2018/file/185e65bc40581880c4f2c82958de8cfe-Paper.pdf>.
- Hung, C. P., Kreiman, G., Poggio, T., & DiCarlo, J. J. (2005). Fast readout of object identity from macaque inferior temporal cortex. *Science*, 310, 863–866.
- Jiang, Z., Chen, T., & Li, M. (2018). Efficient deep learning inference on edge devices. *ACM SysML*.
- Johansson, R. S., & Birznieks, I. (2004). First spikes in ensembles of human tactile afferents code complex spatial fingertip events. *Nature Neuroscience*, 7(2), 170–177.
- Kheradpisheh, S. R., & Masquelier, T. (2020). Temporal backpropagation for spiking neural networks with one spike per neuron. *International Journal of Neural Systems*, 30(06), Article 2050027.
- Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. In *ICLR (poster)*.
- Krizhevsky, A., & Hinton, G. (2009). Learning multiple layers of features from tiny images. *Technical report*, Toronto, Ontario: University of Toronto.
- Kubke, M. F., Massoglia, D. P., & Carr, C. E. (2002). Developmental changes underlying the formation of the specialized time coding circuits in barn owls (*Tyto alba*). *Journal of Neuroscience*, 22(17), 7671–7679.
- Lennie, P. (2003). The cost of cortical computation. *Current Biology*, 13(6), 493–497.
- Mirsadeghi, M., Shalchian, M., Kheradpisheh, S. R., & Masquelier, T. (2021). Stidi-BP: Spike time displacement based error backpropagation in multilayer spiking neural networks. *Neurocomputing*, 427, 131–140.
- Mostafa, H. (2018). Supervised learning based on temporal coding in spiking neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 29(7), 3227–3235. <http://dx.doi.org/10.1109/TNNLS.2017.2726060>.
- Neftci, E. O., Mostafa, H., & 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), 51–63.
- Optican, L. M., & Richmond, B. J. (1987). Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. 3. Information theoretic analysis. *Journal of Neurophysiology*, 57, 162–178.
- Park, S., Kim, S., Na, B., & Yoon, S. (2020). T2FSNN: Deep spiking neural networks with time-to-first-spike coding. In *2020 57th ACM/IEEE design automation conference* (pp. 1–6). IEEE.
- Patterson, D., Gonzalez, J., Le, Q., Liang, C., Munguia, L., Rothchild, D., et al. (2021). Carbon emissions and large neural network training. arXiv preprint arXiv:2104.10350.
- Pillow, J. W., Shlens, J., Paninski, L., Sher, A., Litke, A. M., Chichilnisky, E. J., et al. (2008). Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*, 454, 995–999.
- Rieke, F., Warland, D., de Ruyter van Steveninck, R., & Bialek, W. (1996). *Spikes - Exploring the neural code*. Cambridge, MA: MIT Press.
- Rueckauer, B., & Liu, S. (2018). Conversion of analog to spiking neural networks using sparse temporal coding. In *2018 IEEE international symposium on circuits and systems* (pp. 1–5). IEEE.
- Rueckauer, B., Lungu, I., Hu, Y., Pfeiffer, M., & Liu, S. (2017). Conversion of continuous-valued deep networks to efficient event-driven networks for image classification. *Frontiers in Neuroscience*, 11, 682.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., et al. (2015). ImageNet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3), 211–252. <http://dx.doi.org/10.1007/s11263-015-0816-y>.
- Sebastian, A., Le Gallo, M., Burr, G. W., Kim, S., BrightSky, M., & Eleftheriou, E. (2018). Tutorial: Brain-inspired computing using phase-change memory devices. *Journal of Applied Physics*, 124(11), Article 111101.
- Sorbaro, M., Liu, Q., Bortone, M., & Sheik, S. (2020). Optimizing the energy consumption of spiking neural networks for neuromorphic applications. *Frontiers in Neuroscience*, 14, 662.
- Stanojevic, A., Eleftheriou, E., Cherubini, G., Woźniak, S., Pantazi, A., & Gerstner, W. (2022). Approximating relu networks by single-spike computation. In *2022 IEEE international conference on image processing* (pp. 1901–1905). IEEE.
- Stanojevic, A., Woźniak, S., Bellec, G., Cherubini, G., Pantazi, A., & Gerstner, W. (2023). Are training trajectories of deep single-spike and deep ReLU network equivalent?. <http://arxiv.org/abs/2306.08744>.

- Stockl, C., & Maass, W. (2021). Optimized spiking neurons can classify images with high accuracy through temporal coding with two spikes. *Nature Machine Intelligence*, 3, 230–238.
- Strubell, E., Ganesh, A., & McCallum, A. (2020). Energy and policy considerations for modern deep learning research. In *Proceedings of the AAAI conference on artificial intelligence: Vol. 34*, (09), (pp. 13693–13696).
- Tan, M., & Le, Q. (2019). Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning* (pp. 6105–6114). PMLR.
- Tavanaei, A., Ghodrati, M., Kheradpisheh, S. R., Masquelier, T., & Maida, A. (2019). Deep learning in spiking neural networks. *Neural Networks*, 111, 47–63.
- Thorpe, S., Delorme, A., & Van Rullen, R. (2001). Spike-based strategies for rapid processing. *Neural Networks*, 14, 715–725.
- Thorpe, S., Fize, D., & Marlot, C. (1996). Speed of processing in the human visual system. *Nature*, 381, 520–522.
- Wang, X., Han, Y., Leung, V. C., Niyato, D., Yan, X., & Chen, X. (2020). Convergence of edge computing and deep learning: A comprehensive survey. *IEEE Communications Surveys & Tutorials*, 22(2), 869–904.
- Widmer, S., Kossel, M., Cherubini, G., Woźniak, S., Francese, P. A., Stanojevic, A., et al. (2023). Design of time-encoded spiking neural networks in 7nm CMOS technology. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 1. <http://dx.doi.org/10.1109/TCSII.2023.3277784>.
- Woźniak, S., Pantazi, A., Bohnstingl, T., & Eleftheriou, E. (2020). Deep learning incorporating biologically inspired neural dynamics and in-memory computing. *Nature Machine Intelligence*, 2(6), 325–336.
- Xiao, H., Rasul, K., & Vollgraf, R. (2017). Fashion-mnist: A novel image dataset for benchmarking machine learning algorithms. arXiv preprint [arXiv:1708.07747](https://arxiv.org/abs/1708.07747).
- Yamins, D. L. K., & DiCarlo, J. J. (2016). Using goal-driven deep learning models to understand sensory cortex. *Nature Neuroscience*, 19, 356–365.
- Yan, Z., Zhou, J., & Wong, W. (2021). Near lossless transfer learning for spiking neural networks. In *Proceedings of the AAAI conference on artificial intelligence: Vol. 35*, (12), (pp. 10577–10584).
- Yang, K., Yau, J. H., Deng, J., & Russakovsky, O. (2022). A study of face obfuscation in ImageNet. In *Proc. 39th intern. conf. machine learning: Vol. 162*, (1), (pp. 25313–25330). PMLR.
- Yang, K., Yau, J. H., Fei-Fei, L., Deng, J., & Russakovsky, O. (2022). A study of face obfuscation in ImageNet. In K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, & S. Sabato (Eds.), *Proceedings of machine learning research: Vol. 162, Proceedings of the 39th international conference on machine learning* (pp. 25313–25330). PMLR, URL <https://proceedings.mlr.press/v162/yang22q.html>.
- Zenke, F., & Ganguli, S. (2018). Superspike: Supervised learning in multilayer spiking neural networks. *Neural Computation*, 30(6), 1514–1541.
- Zenke, F., & Vogels, T. P. (2021). The remarkable robustness of surrogate gradient learning for instilling complex function in spiking neural networks. *Neural Computation*, 899–925.
- Zhang, M., Wang, J., Wu, J., Belatreche, A., Amornpaisannon, B., Zhang, Z., et al. (2021). Rectified linear postsynaptic potential function for backpropagation in deep spiking neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 33(5), 1947–1958.
- Zhang, L., Zhou, S., Zhi, T., Du, Z., & Chen, Y. (2019). Tdsnn: From deep neural networks to deep spike neural networks with temporal-coding. In *Proceedings of the AAAI conference on artificial intelligence: Vol. 33*, (01), (pp. 1319–1326).
- Zhou, B. (2018). Github. URL <https://github.com/zhoubolei>.
- Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., & Torralba, A. (2017). Places: A 10 million image database for scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(6), 1452–1464.