NEUROMORPHIC
Computing and Engineering

Simple and complex spiking neurons: perspectives and analysis in a simple STDP scenario

Davide L Manna[1,*] ⓘ, Alex Vicente-Sola[1] ⓘ, Paul Kirkland[1] ⓘ, Trevor J Bihl[2] ⓘ
and Gaetano Di Caterina[1] ⓘ

1   Neuromorphic Sensor Signal Processing Lab, Centre for Image and Signal Processing, Electrical and Electronic Engineering,
    University of Strathclyde, Glasgow, United Kingdom
2   Air Force Research Laboratory, Wright Patterson AFB, OH, United States of America
*   Author to whom any correspondence should be addressed.

E-mail: davide.manna@strath.ac.uk

Abstract

Spiking neural networks (SNNs) are largely inspired by biology and neuroscience and leverage ideas and theories to create fast and efficient learning systems. Spiking neuron models are adopted as core processing units in neuromorphic systems because they enable event-based processing. Among many neuron models, the integrate-and-fire (I & F) models are often adopted, with the simple leaky I & F (LIF) being the most used. The reason for adopting such models is their efficiency and/or biological plausibility. Nevertheless, rigorous justification for adopting LIF over other neuron models for use in artificial learning systems has not yet been studied. This work considers various neuron models in the literature and then selects computational neuron models that are single-variable, efficient, and display different types of complexities. From this selection, we make a comparative study of three simple I & F neuron models, namely the LIF, the quadratic I & F (QIF) and the exponential I & F (EIF), to understand whether the use of more complex models increases the performance of the system and whether the choice of a neuron model can be directed by the task to be completed. Neuron models are tested within an SNN trained with spike-timing dependent plasticity (STDP) on a classification task on the N-MNIST and DVS gestures datasets. Experimental results reveal that more complex neurons manifest the same ability as simpler ones to achieve high levels of accuracy on a simple dataset (N-MNIST), albeit requiring comparably more hyper-parameter tuning. However, when the data possess richer spatio-temporal features, the QIF and EIF neuron models steadily achieve better results. This suggests that accurately selecting the model based on the richness of the feature spectrum of the data could improve the whole system's performance. Finally, the code implementing the spiking neurons in the SpykeTorch framework is made publicly available.

1. Introduction

As technology in the neuromorphic (NM) computing field keeps on advancing, so are the software methodologies and algorithms that can leverage the low-power, low-latency and event-driven properties that characterize NM [1]. Often, inspired by the success of conventional deep learning, this results in the development of spiking neural networks (SNNs). When it comes to designing an SNN learning system for some machine learning task, researchers are faced with many decisions to make. Among these comes the choice of a particular neuron model. This specific aspect of the development of an SNN is an extremely sensitive one as spiking neurons are the core processing units of an NM system. To draw a parallel with the conventional deep learning (DL) research, spiking neurons can be thought of as being activation functions (such as the ReLU, ATAN etc), but holding an internal state. The dynamics of this state through time are governed by the differential equations that constitute the spiking neuron model.

Different neuron models exhibit different state dynamics. From a neuroscience point of view, these differences are immediately clear [2]. Some models are able to capture certain intrinsic behaviours of neurons, e.g. they can burst, chatter or fast-spike, while others cannot. Some models are also better at approximating subthreshold dynamics, thus possibly being more accurate representations of real neurons. However, it is still unclear how this ability translates into applicability in SNNs. There is in fact no definite answer onto whether certain types of neural dynamics can be beneficial to particular SNN applications, nor any common knowledge on the criteria that should drive the choice of such neurons in relation to such dynamics.

Within some specific contexts, the choice is constrained by the available hardware. As a matter of fact, several neuromorphic chips allow to only adopt the specific neuron model that the chip is able to emulate. BrainScaleS [3] for instance allows to only adopt the adaptive exponential integrate-and-fire neuron model; neuroGrid [4] allows only an adaptive quadratic integrate-and-fire model; Loihi [5], SyNAPSE [6, 7] and TrueNorth [8] allow only leaky integrate-and-fire (LIF) based models; BiCoSS [9] allows using different models, including the LIF, Izhikevich's, and Hodgkin–Huxley's models, for large-scale SNNs and brain simulations. Until recently, the only chip allowing the implementation of any type of neuron model was SpiNNaker [10]. However, the recently released Loihi 2 adds to the list of chips with programmable neuron models [11], hence highlighting the importance of an accurate investigation on this matter.

It is thus interesting to look at what are the most suitable neurons models for SNN development. This can help to understand if the dynamics of the neurons relate, in some way, to the dynamics of the spatio-temporal features of the data.

Simple LIF neurons are the de-facto standard choice when it comes to SNN design [12–24]. When a rationale for this is provided, this choice is often attributed to the simplicity and, consequently, to the efficiency of the LIF neuron model. Whatever the case, such reasons hardly account for the accuracy performance of the task at hand, neither for the temporal feature representation capability of the model. Other works rely on more complex neuron models [25–28], attributing the choice to the biological plausibility or, once again, to efficiency. Neurons with different dynamics are present in areas of the brain with different functionalities [29–31]; however, the same does not apply to spiking neuron models in ML, where often the simplest neurons are used, leaving it unclear whether there are advantages or disadvantages relative to different types of NM data.

This work aims to answer the following research questions:

- Does the chosen neuron model influence the performance of an SNN?
- Should the choice of the neuron model be related with the data it will have to process?

Specifically, we are interested in understanding whether neuron models with different and more complex neuronal dynamics display any advantages over simpler (LIF) ones in an unsupervised learning context. Furthermore, we investigate whether such differences might exist depending on the task set to the network; hence we perform experiments using two different datasets. To do this, we first develop a basic experiment, which represents a simple yet efficient way to start a comparison between neuron models. We select neuron models so that they scale up in terms of complexity and spiking patterns, and evaluate their performance within the same neural network architectures on the N-MNIST dataset [32]. Then, we use the same neural network and train it on classification tasks taken from the DVS gestures dataset [33], which displays a richer distribution of events. In order to perform the aforementioned experiments, we further contribute by enriching the SpykeTorch [34] framework with a new set of spiking neurons[3].

The rest of this paper is organized as follows: in section 2 we provide some background information regarding the multitude of neuron models found in the literature and highlight some relevant related works; in section 3 we present our experimentation pipeline in detail, focusing on the datasets, neural network design and learning paradigms; sections 4 and 5 contain respectively the results obtained through our experiments and initiate an in-depth discussion on such results; section 6 concludes the paper.

2. Background and related works

Spiking neuron models were first born in the field of neuroscience and neurophysiology, where mathematical models were developed to reproduce what was found by recording the activity of real neurons [35–37]. Early SNNs resulted from studies aimed at understanding biological dynamics [38], or simulating areas of the brain and neuronal interactions [39–41]. The shift towards their use for computational tasks was gradual and was formalised afterwards [42, 43], with the focus generally being on the overall computational abilities of the network. When it comes to developing SNNs for ML applications, spiking neurons can be thought of as stateful

---

[3] Code available at https://github.com/daevem/SpykeTorch-Extended.

activation functions. This means they retain a state of their value (the membrane potential) reached through previous inputs. They are thresholding functions, therefore allowing to only forward information upon the reaching of a set threshold.

In conventional DL, activation functions have been extensively studied due to their importance in the propagation of the information. Nonlinear functions such as the sigmoid and Tanh were introduced to break the linearity of multilayer perceptrons [44, 45]. rectified linear units (ReLUs) substituted them to solve the vanishing gradient problem and allow deeper networks. Further variants [46–49] addressed other issues like the dying ReLU problem and helped to improve the performance of the networks [50–55]. Instead, in the context of SNNs and spiking neurons, it is hard to find works in the literature relative to the differences in the use of different neuron models in SNNs for NM and DL applications. To the best of our knowledge, the only work considering the role of spiking neurons in learning from a DL point of view is the one by Traub *et al* [56]; however, they focus on the qualitative properties of spike-timing-dependent plasticity (STDP) and the mean firing rate of the system after training. Furthermore, they consider a non-NM dataset and a different set of neurons.

Most of the works on spiking neurons concentrate on the neurobiological aspects they expose. One of the most influential works in this matter is the one by Izhikevich [57], which compared several models of spiking neurons (LIF, LIF with adaptation, LIF-or-burst, resonate-and-fire, QIF, Izhikevich's, FitzHugh–Nagumo, Hindmarsh–Rose, Morris–Lecar, Wilson, Hodgkin–Huxley), outlining their ability to reproduce observed neuronal behaviours and the cost (in terms of floating-point operations) of implementing such neurons in software applications. Similar work was conducted in [58], but focusing on a smaller subset of neurons (LIF, Izhikevich's, FitzHugh–Nagumo, Wilson, Hodgkin–Huxley) and analysing their numerical stability. Although they closely study spiking neuron models, the two studies above concentrate on their computational costs and the intrinsic biological mechanics that each model can reproduce. However, they do not consider the effect of using spiking neurons with different dynamics in a DL system.

A number of other works concentrate on the efficacy of neuron models in representing observed cortical neurons firing patterns. One example is given by [59], where the authors make an exploratory analysis of how parameters influence Izhikevich's neurons in showing different spiking patters. Still regarding Izhikevich's neurons, Kumar *et al* [60] estimate parameters that allow the neuron model to optimally reproduce a given spike train. Teeter *et al* [61] use a generalized version of the LIF neuron model (GLIF) to understand whether more complex models allow to predict spike timing behaviors more closely; they conclude that this ability does not increase monotonically with the complexity, nor with the ability to reproduce sub-threshold dynamics. In [62] it is argued that integrate-and-fire (I & F) neuron models are good enough estimators of input spike trains when coupled with an adaptation variable. This is both quantitatively and qualitatively shown by the authors and provides a good ground to the adoption of this family of neuron models.

Finally, in [63] the authors compare different neuron models embedded in a liquid state machine (LSM), a particular type of reservoir computing network. They evaluate their model with two different input patterns and use Euclidean distance and entropy to estimate the 'separation' ability of the LSM, i.e. its ability to generate different response patterns for different stimuli. They test their LSM using six spiking neuron models (I & F, resonate-and-fire, FitzHugh–Nagumo, Hindmarsh–Rose, Morris–Lecar, and Izhikevich's) and perform experiments varying the density of the connections between the neurons. They found that the LSM failed to achieve satisfactory levels of separation only when using Izhikevich's neurons. Other models allowed better separation levels depending on the density of the connections. They conclude by postulating that, for LSM implementations, Morris–Lecar, resonate-and-fire, and Hindmarsh–Rose models are most suitable.

2.1. Neuron models in the literature

Neurons in the human brain differ by several characteristics, ranging from the type of neurotransmitters used, to the shape they have and the spiking patterns they expose. As a common ground, however, they all share a basic structure composed of a dendritic tree (input channels), an axon (output mean), and a soma (core) [2]. Dendrites (and axons) can be further broken down into several blocks through which signals travel. In the literature, the neurons have been modelled with different levels of abstraction and detail. Works such as [64–68] focus on the definition and use of so-called multi-compartment neuron models. These kind of models try to account for each compartment of a neuron (dendrites, axon and soma) individually to more closely replicate biological evidence. However, a large number of models considers the neurons as dimensionless entities (point neurons), hence focussing on the modelization of the soma only. They can be roughly subdivided in two larger groups (see figure 1), the bio-physical or conductance-based models and the event-based or integrate-and-fire models. In this work, we concentrate on the study of point neurons as they are largely used in SNNs for ML applications.

Figure 1. Venn diagram of some spiking neurons.
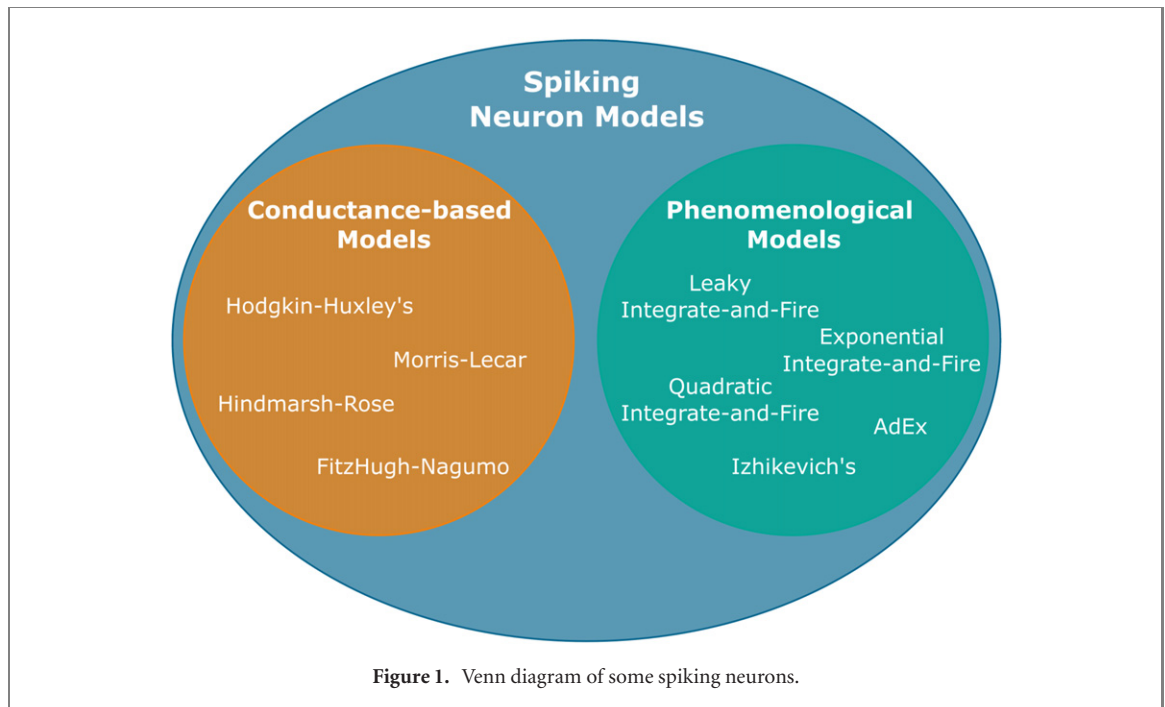
2.1.1. Conductance-based models
This class of neuron models is characterized by the fact that all the variables and parameters present in the model have a biophysical correspondence and are therefore measurable through experiments [69]. Among them, the Hodgkin–Huxley (HH) model is considered to be one of the most important in computational neuroscience and defines a system of four non-linear differential equation with four variables and a number of parameters. While further levels of complexity can be attained by including further variables in the model, this is not amenable to mathematical analysis. In fact, other simpler conductance-based models have been derived in the literature in order to ease the analysis, while still retaining biophysical plausibility. Some examples are the FitzHugh–Nagumo model [70, 71], the Hindmarsh–Rose [72] and the Morris–Lecar model [73]. Nevertheless, they still remain rather complex for what concerns analysis and computation, therefore this family of neuron models is often used only when studying single-cell or small population dynamics [57].

2.1.2. Phenomenological models
The family of integrate-and-fire or phenomenological neuron models comprises all those models that treat spikes as stereotypical events in time [2]. Therefore, each spike is completely described by the time at which it occurred, or was emitted. Integrate-and-fire models require at least two equations, one describing the dynamics of the membrane potential and the other one defining the action potential generation. Events are integrated over time and convey electrical charges that can cause excitation or inhibition of the membrane potential of the receiving neuron. Differently from the conductance-based models, the phenomenological ones are more indicated for the development of neural networks [57, 74, 75], thanks to their overall lower complexity and the lower number of parameters, which enable easier fitting.

The simplest model, apart from the perfect integrator, is the LIF [35]. The dynamics of the membrane potential are here described by the following linear differential equation:

$$\tau_{\mathrm{m}}\frac{\mathrm{d}u}{\mathrm{d}t} = -(u(t) - u_{\mathrm{rest}}) + R \cdot I, \tag{1}$$

where τ_{m} is the membrane time constant, $u(t)$ is the membrane potential as a function of time, u_{rest} is the resting potential of the membrane, R is a resistance and I is the incoming current. Here, the term $-(u(t) - u_{\mathrm{rest}})$ accounts for the leakage of the membrane potential.

Although this model lacks the ability to describe most of the neuronal dynamics, it is the most common choice for the development of large scale neural networks, mostly because of its efficiency.

More complex I & F models attempt to account for some non-linear dynamics of neurons as a function of the value of their membrane potential in a certain moment in time. Two examples are given by the exponential integrate-and-fire (EIF) model [76] and by the quadratic integrate-and-fire neuron (QIF) or theta neuron

[69, 77]. As shown in equation (2), EIF model expands on the LIF model by including an exponential dependency on the current state of the membrane potential:

$$\tau_m \frac{\mathrm{d}u}{\mathrm{d}t} = -(u(t) - u_{\text{rest}}) + \Delta_T \, \exp\left(\frac{u(t) - \Theta_{\text{rh}}}{\Delta_T}\right) + R \cdot I, \tag{2}$$

where Δ_T is a parameter determining the sharpness of the exponential curve and Θ_{rh} is the rheobase threshold. When $u > \Theta_{\text{rh}}$, the exponential term becomes prominent over the linear one, leading to an upswing of the curve that takes the membrane potential to infinity in finite time.

The QIF model, given by equation (3), employs a quadratic dependency from the membrane potential:

$$\tau_m \frac{\mathrm{d}u}{\mathrm{d}t} = a_0(u(t) - u_c)(u(t) - u_{\text{rest}}) + R \cdot I, \tag{3}$$

where a_0 is a parameter of the model that regulates the magnitude of the dependency from the membrane potential and u_c is a cut-off threshold such that, when $I = 0$ and $u > u_c$, the membrane potential grows until the emission of a spike.

Both the QIF and the EIF bring in a further level of complexity with the inclusion of non-linear dependencies that affect both the computational costs and the ease of analysis, but allow for a more precise generation of spikes [2]. Additionally, a hidden cost lies in the use of two extra parameters in each of them.

2.1.3. Other multi-variable I & F models

With the inclusion of adaptation variables within the neuron model, it is possible to account for a larger number of spiking patterns and to render possible the manifestation of spike bursts, spike-adaptation responses and irregular spiking [2]. This comes at the cost of more differential equations in the model (one per variable) and two relevant examples are given by the adaptive exponential integrate-and-fire (AdEx) [78] neuron model, which builds on top of the EIF, and by Izhikevich's neuron model [79] which builds on top of the QIF.

A number of neuron models have been theorized in the literature, all answering to different modelling needs or considering different aspects of the observed neuronal behaviours. The ones cited above are amongst the most relevant for what concerns this study and NM computing. In fact, as reported above, the LIF model is the most widely used in the development of SNN for NM applications, but at the same time, models like the AdEx and Izhikevich's have received a lot of attention in the literature. The EIF and the QIF are on one hand the baseline of the AdEx and Izhikevich's models respectively, and, on the other hand, a slightly more complex single-variable alternative to the LIF neuron model. As such, in this study, we will focus on these single-variable models.

3. Methods

We are interested in assessing the performance of different neuron models within the context of a spiking convolutional neural network (SCNN) trained with STDP. Since many factors could determine the outcome of the training, we begin by designing a simple experiment which involves the minor number of structural elements possible. This is done in order to limit the number of components that might impact the overall system performance. Therefore, we use a single-layer convolutional network in which spiking neurons are embedded right after the convolution operation on the input. The task set to the SCNN is a binary classification task, with the pairs of classes taken from the neuromorphic MNIST dataset [32] and the DVS Gestures dataset [33], which contain event-based data samples. To develop the learning pipeline, we utilize SpykeTorch [34] as a base framework and build on top of it to include the elements required by this study, such as the diverse spiking neuron models.

3.1. Event-based data

To assess the performance of our simple network, we select the two natively neuromorphic datasets mentioned above. We purposely discard other non-native NM datasets as they do not possess a temporal domain, nor data is originally event-based.

Data in the N-MNIST dataset is collected by recording MNIST digits shown on a screen using a moving DVS camera. Specifically, the camera makes the same three predefined movements for every sample, each lasting roughly 100 ms. In this way, although the dataset is built on top of a non-neuromorphic one, data samples in the dataset are natively event-based, rather than being converted from a static image. Figure 2 reports the count of events per time for two example classes and throughout all the classes of the dataset. By contrast, data samples in the DVS gestures dataset (see figure 3 for the inter-class distribution of events over time) are recorded using a fixed DVS camera in front of which participants move their arms according to
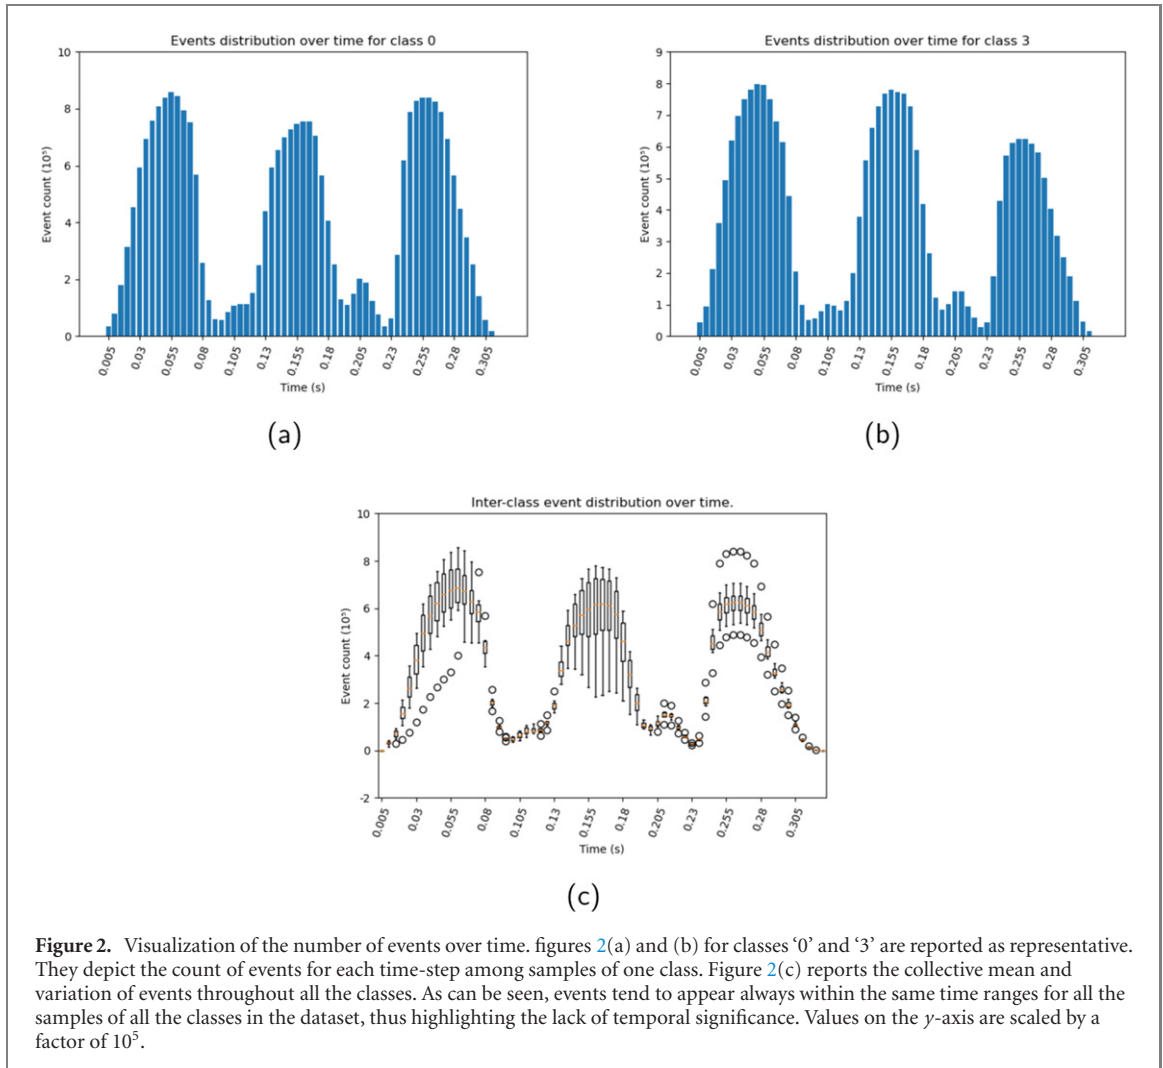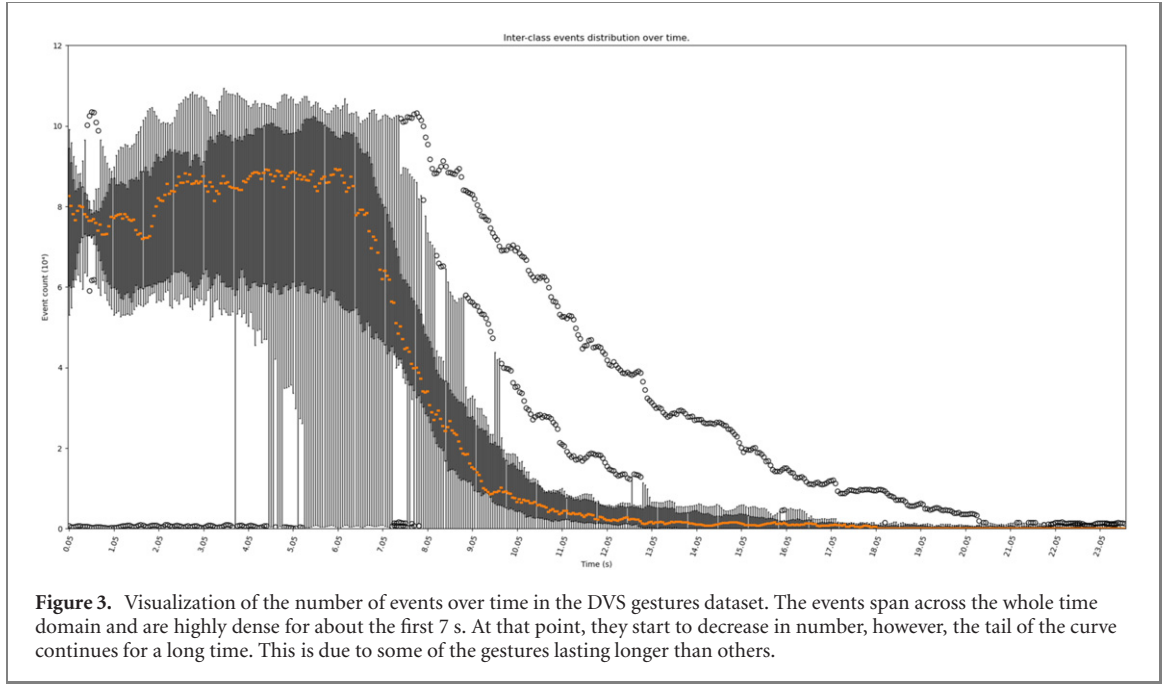
Figure 2. Visualization of the number of events over time. figures 2(a) and (b) for classes '0' and '3' are reported as representative. They depict the count of events for each time-step among samples of one class. Figure 2(c) reports the collective mean and variation of events throughout all the classes. As can be seen, events tend to appear always within the same time ranges for all the samples of all the classes in the dataset, thus highlighting the lack of temporal significance. Values on the *y*-axis are scaled by a factor of 10^5.

instructions. Thus, 11 different classes of gestures are obtained, including for example arm rotation, waving or performing air guitar. The 11th class encodes 'other' random movements and is not considered in this work for simplicity.

Event data comes in the form of address event representation (AER)-encoded files in which every sample is constituted by a sequence of events. Events are characterized by the specific time at which they occurred, by the location on the 2D plane and by the polarity (negative or positive light change). Similarly to [80], to make data usable by a 2D convolutional neural network (CNN) we populate a 2D image using all the events that took place between time t and $t + \mathrm{d}t$, allowing at most 1 event per (x, y) coordinates. For simplicity, we consider all events as being positive and use a batch size of 1. As a result, the network processes only 1 event map with a time resolution $\mathrm{d}t$ belonging to only 1 sample at a time. Finally, we characterize each event with a value of $\frac{1}{t_s}$ in line with [2]. This is done in order to preserve the amount of charge that a spike carries regardless of the time-step (t_s) size. No further pre-processing is applied to the data.

3.2. Spiking neurons implementation

The phenomenological family of neuron models is the best option when developing SNNs, as outlined in section 2.1. We specifically concentrate on three integrate and fire neurons, namely the LIF, the QIF and the EIF. The parameters used for these models in the same-parameter setting of experiments (see section 4) are reported in table 1. The LIF is the most widely used neuron that embeds a time dependency through the membrane potential leakage. The QIF and the EIF represent valid alternatives given their ability to best fit observed cortical neurons [2, 69]. Since they are single-variable models, they stand for a fairer comparison with the LIF, which is single-variable too. Indeed they all depend on the value of the membrane potential, but they all employ different types of dependencies from it. Furthermore, they are the base on which other more complex and popular neuron models are built on, respectively Izhikevich's neuron and the AdEx neuron model.

Figure 3. Visualization of the number of events over time in the DVS gestures dataset. The events span across the whole time domain and are highly dense for about the first 7 s. At that point, they start to decrease in number, however, the tail of the curve continues for a long time. This is due to some of the gestures lasting longer than others.
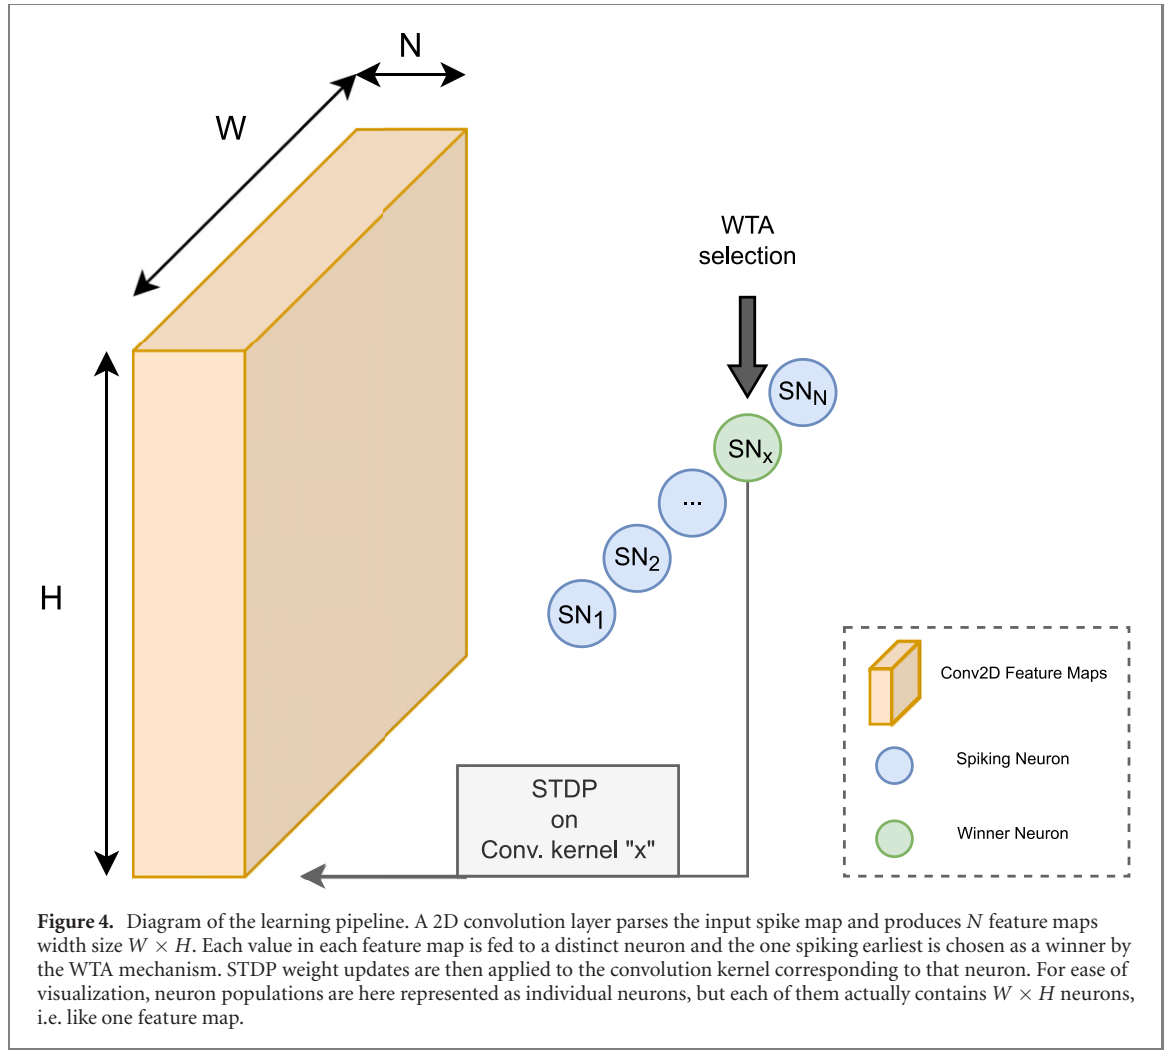
Table 1. Table of hand-tuned parameters used for neurons. Each column represents a different parameter, as outlined in section 2.1. The capacitance C is used instead of the resistance R by leveraging the equality $R = \frac{\tau_{\mathrm{m}}}{C}$. For every neuron, the time-step size used was 0.02, and the voltage threshold was recalculated every 100 samples.

Neurons	τ_{m}	u_{rest}	C	Δ_T	Θ_{rh}	a	u_{c}
LIF	0.2	0	0.1	—	—	—	—
EIF	0.2	0	0.1	1352	216	—	—
QIF	0.2	0	0.1	—	—	0.01	216

The SpykeTorch framework comes with a simple version of a I & F neuron model. To enable the use of these neuron models for our experiments, we expand on the framework and implement the above models by adapting the equations provided in [2]. State updates in the neurons are thus calculated on a per time-step basis, where each call to the neuron layer corresponds to an advancement of t_{s} time from the previously calculated update. Each neuron layer generates a number of neurons that reflects the size of the incoming post-synaptic currents, multiplied by the number of neuron populations that was specified at creation time. This allows for a more seamless inclusion in any point of an SNN. When a neuron in a population emits a spike, all the other neurons belonging to the same population are inhibited and put in a refractory state to promote learning in other populations. Each neuron model is implemented as a class inheriting from a parent neuron class, in an object-oriented programming style. This differs from the original SpykeTorch implementation style, however, this approach was required for neurons that maintain an internal state. Besides, compatibility with the modules and neurons in SpykeTorch is maintained. Further features and details are present in the actual implementation, however, these are not relevant to the current study and the authors point the reader to [81] and to the dedicated repository page for in-depth descriptions.

3.3. SCNN and learning

We develop a simple feedforward convolutional network for our experiments, with a single convolutional layer that parses the input and connects it to the spiking neurons. Figure 4 illustrates the pipeline for a better understanding. The weights of the kernel can be thought of as being synapse strengths and the resulting feature map is the input to a set of spiking neurons arranged accordingly. In other words, the convolution represents the connectivity scheme for the spiking neurons. The use of a convolutional layer to connect inputs with the spiking neurons allows them to more easily learn spatial features. The weights of the convolutional layer are the only parameters being learnt by the network and no pooling nor normalization is applied.

Figure 4. Diagram of the learning pipeline. A 2D convolution layer parses the input spike map and produces N feature maps width size $W \times H$. Each value in each feature map is fed to a distinct neuron and the one spiking earliest is chosen as a winner by the WTA mechanism. STDP weight updates are then applied to the convolution kernel corresponding to that neuron. For ease of visualization, neuron populations are here represented as individual neurons, but each of them actually contains $W \times H$ neurons, i.e. like one feature map.

We use the STDP presented and used in [82–84] as an unsupervised learning rule. By employing this kind of learning rule and, therefore, using it to evaluate neurons, we want to take a step closer to local feedforward learning, which is believed to be the key to exploit the potentials of neuromorphic engineering at its best [1, 85]. The mentioned STDP rule applies the following weight updates:

$$\Delta W_{i,j} = \begin{cases} A^+ \times (W_{i,j} - \text{LB}) \times (\text{UB} - W_{i,j}) & \text{if } T_j \leqslant T_i, \\ A^- \times (W_{i,j} - \text{LB}) \times (\text{UB} - W_{i,j}) & \text{if } T_j > T_i, \end{cases} \qquad (4)$$

where $W_{i,j}$ is the weight of the synapse connecting neuron j (pre-synaptic) to neuron i (post-synaptic), LB and UB are a lower and an upper bound value respectively, T_j is the timing of the spike emitted by neuron j, T_i the timing of the spike emitted by neuron i, and A^+ and A^- are two parameters used to scale the weight update.

It is thus a system of two parabolic equations that are applied depending on whether long term potentiation (LTP) or long term depression (LTD) needs to take place. One of the consequences of using this learning rule is that weight updates are self-regularized. In fact, the closer the weights get to the boundary values, the smaller the updates will be, allowing the weights to be refined more granularly as the learning proceeds. Another aspect that is important to highlight is in how this learning rule is applied. While the original theorization of Hebbian rules such as STDP states that the weight update should be proportional in value and sign on the time-difference between the post- and the pre-synaptic spikes, in a software implementation some approximations are required. Therefore, for each time-step of the execution, we apply LTP on weights connected to input locations where there has been a spike, and LTD in all the others. In order to promote competitive and differentiated feature learning, we also employ a k-winners take all (WTA) (with $k = 1$) learning paradigm. WTA allows only k neurons per time-step to be eligible for STDP updates, specifically the ones firing sooner.

Finally, as an homeostatic mechanism to allow neurons to keep on firing despite the changes in their synaptic weights, we re-calculate individual neurons' thresholds according to equation (5). This for of adaptive

Table 2. Table of optimized parameters used for neurons. Each column represents a different parameter, as outlined in section 2.1. The capacitance C is used instead of the resistance R by leveraging the equality $R = \frac{\tau_m}{C}$. For every neuron, the time-step size used was 0.02, and the voltage threshold was recalculated every 100 samples. HC stands for hand clapping, while RHW stands for right hand wave.

Neurons	0 vs 1							HC vs RHW						
	τ_m	u_{rest}	C	Δ_T	Θ_{rh}	a	u_c	τ_m	u_{rest}	C	Δ_T	Θ_{rh}	a	u_c
LIF	0.0602	0	0.2983	—	—	—	—	0.1435	0	0.3020	—	—	—	—
EIF	0.2578	0	0.2178	32	91.14	—	—	0.2389	0	0.2086	32	69.42	—	—
QIF	0.2804	0	0.2178	—	—	0.001	69.72	0.0621	0	0.1026	—	—	0.0393	275.95

thresholding is often required when employing STDP, and a common practice for SNNs [83].

$$V_{\text{thresh}} = \lambda \cdot R \cdot A \cdot \frac{t_s}{\tau_m} \cdot \overline{W} \cdot (W_k \cdot H_k) \cdot n_c, \tag{5}$$

and since $\tau_m = R \cdot C$, we can equivalently re-write:

$$V_{\text{thresh}} = \lambda \cdot A \cdot \frac{t_s}{C} \cdot \overline{W} \cdot (W_k \cdot H_k \cdot N_c), \tag{6}$$

where A is the amplitude of the spike, in our case assigned to be $A = \frac{1}{t_s}$, R is the resistance, C is the capacity, τ_m is the membrane time constant, $\overline{W}$ is the average value of the synaptic weights, W_k, H_k and N_c are the width, height and depth (number of channels) of the synaptic kernel, and λ is a regulation parameter that takes values in the range $[0, 1]$. In equation (6), the term $A \cdot \frac{t_s}{C} \cdot \overline{W}$, can be explained as being the average effect perceived on the membrane potential as a result of a single spike, whereas the second term, $(W_k \cdot H_k) \cdot n_c$, scales this effect to the size of the synaptic kernel. Therefore, equation (6) calculates what would be the average post-synaptic potential perceived in the case the input was dense with spikes. The parameter λ serves as a regulation of what percentage of this amount would be necessary to reach before emitting a spike.

3.4. Classification

Because we employ an unsupervised learning rule, labels are not used at any point during the learning of the weights. However, labels are needed to classify data samples, therefore we adopt a system similar to [86, 87] and count, for each neuron, the number of times it spiked in response to samples having a given label. At the end of the training phase, each neuron is assigned the label for which it spiked the most during training.

During the inference phase, for each data sample, a sequence of spikes is collected and weighted depending on the order they arrived. These are then summed and the label corresponding to the highest value is selected as the classification outcome.

3.5. Hyper-parameter optimization

Defining a good set of parameters for a machine learning system is a non-trivial task. This often requires a lot of expertise and hand tuning and is greatly error-prone. To reduce the possibility of selecting sub-optimal parameters for neurons which would result in poor performance, we thus make use of an optimization system to find reasonably good parameters for our experiments. To this extent, we employ the BOHB optimization [88] using the HpBandSter library. This technique combines Bayesian optimization (BO) and hyperband (HB), a resource allocation and early stopping strategy. To use BOHB, we adapted our implementation so that it could be optimized using the HpBandSter library. More importantly, we defined the domains in which every parameter was allowed to vary. For example, the time-constant τ_m could be drawn from the interval $[0.06, 0.26]$. Moreover, because the optimization process could be task-specific, we performed separate optimizations for each different task. We summarize the final parameters in table 2. As a result of this process, we can draw more robust conclusions about the performance of the neurons.

4. Results

We train the SCNN outlined in section 3 using a subset of the N-MNIST and the DVS gestures datasets. We randomly select four distinct couples of classes from each dataset and define, for each, a separate binary classification task. In this way, we aim to obtain more generalizable results and to reduce the dependency of the results on a particular coupling. Furthermore, the order in which data is presented to the SCNN might be influential in a system employing STDP as a learning rule. This is due to the fact that STDP rewards and builds on the inputs that are presented earlier. Therefore, to ensure independence from this behaviour of STDP, we repeat every experiment a total of 11 times per task.

Table 3. Table of results on the N-MNIST dataset. In each cell, the mean accuracy ± the standard deviation values are followed by the best accuracy found (after the comma). Values are rounded up to the closest second decimal value.

Neuron model	0 vs 1	2 vs 9	3 vs 7	4 vs 8
LIF	0.77 ± 0.11, 0.93	0.74 ± 0.06, 0.79	0.73 ± 0.04, 0.78	0.57 ± 0.02, 0.59
EIF	0.80 ± 0.12, 0.94	0.64 ± 0.07, 0.71	0.65 ± 0.04, 0.71	0.55 ± 0.04, 0.61
QIF	0.57 ± 0.03, 0.61	0.54 ± 0.03, 0.58	0.51 ± 0.02, 0.55	0.52 ± 0.02, 0.55

Table 4. Table of results on the DVS gestures dataset. In each cell, the mean accuracy ± the standard deviation values are followed by the best accuracy found (after the comma). Values are rounded up to the closest second decimal value. In the table, HC stands for hand clapping, RHW for right hand wave, RACW for right arm clockwise, AG for air guitar, RACCW for right arm counter clockwise, AR for arm roll, LACW for left arm clockwise and AD for air drums.

Neuron model	HC vs RHW	RACW vs AG	RHCW vs AR	LACW vs AD
LIF	0.53 ± 0.06, 0.60	0.50 ± 0.05, 0.56	0.54 ± 0.13, 0.66	0.52 ± 0.09, 0.69
EIF	0.58 ± 0.12, 0.77	0.50 ± 0.04, 0.58	0.53 ± 0.11, 0.66	0.46 ± 0.05, 0.52
QIF	0.57 ± 0.10, 0.71	0.48 ± 0.04, 0.52	0.62 ± 0.06, 0.67	0.53 ± 0.09, 0.75

4.1. Same hyper-parameters training

We first conduct our experiments using hand-tuned hyper-parameters. These were found by a trial-and-error practice and represent a set of parameters that enabled learning for the task at hand. This means that neurons using these parameters were able to emit spikes and to have the weights adjusted in a way that enabled the learning of representations of the inputs. Where possible, we adopt the same hyper-parameters for all the neurons in all the experiments on each dataset. Since the QIF and the EIF models introduce two different hyper-parameters each, each of them undergoes a further hand-tuning. Results of the training sessions are shown in table 3 for the N-MNIST-based tasks and in table 4 for the DVS gesture-based tasks. Here, for every task and neuron model, are reported the average and best test accuracies achieved, calculated according to equation (7):

$$accuracy = \frac{TP + TN}{TP + TN + FT + FN},$$

(7)

where TP stands for true positive, TN for true negative, FT for false true and FN for false negative. On each column (task), the best average score and the absolute best score are highlighted in bold.

By examining the results on the N-MNIST-based tasks in table 3, the LIF neuron model is found to perform better than the other two on average. Indeed, the EIF has higher average accuracy only on the 0 vs 1 task, whereas the QIF model fails to achieve accuracy levels high enough to match any of the other two counterparts.

Considering the results reported in table 4, the situation differs slightly. The performance of both the LIF and the EIF neuron models, on average, decreases drastically, whereas the QIF maintains similar levels of accuracy as in the N-MNIST case. Nevertheless, both the EIF and QIF demonstrate superior classification abilities throughout and their top accuracy levels often surpass those of the LIF model. These trends in the accuracy levels highlight two main aspects. Firstly, the complexity of the N-MNIST data is lower than that of the DVS Gestures dataset. Indeed, in both cases, the same neural network architecture and neuron models were employed, yet the accuracy in the DVS gestures is on average considerably lower, hence highlighting the greater difficulty of the task. Data samples in the DVS gestures dataset arguably have richer visual and temporal features which render the learning more difficult when compared to the N-MNIST. Secondly, the richer temporal diversity of the features might be better represented by means of neurons with richer voltage dynamics, such as the QIF and EIF. As reported by the experiments, in fact, these two are steadily better than the LIF models and even though in some instances one performs more poorly, the other still attains higher accuracy, possibly as a result of a better affinity to the temporal dynamics found in that particular task.

4.2. Optimized hyper-parameters training

As a second set of experiments, we employ the optimization system outlined in section 3.5 to obtain a set of hyper-parameters that is heuristically optimal for a specific scenario. The optimization is carried out on the '0 vs 1' and on the 'HC vs RHW' tasks for each neuron model individually. We allow a total of 24 optimization iterations for each neuron and task, to not favor any experiment over the others. Results are reported in table 5. Once again, after obtaining the optimized hyper-parameters, we train and evaluate each model a total of 11 times to increase the robustness of the results. Since the optimization is task-specific, we only evaluate our models on the two representative tasks they were optimized on. In the case of the '0 vs 1' task based on

Table 5. Table of results using optimized hyper-parameters. In each cell, the mean accuracy ± the standard deviation values are followed by the best accuracy found (after the comma). Values are rounded up to the closest second decimal value. Optimization and evaluation is performed on one representative task per dataset only.

Neuron model	0 vs 1	HC vs RHW
LIF	$0.95 \pm 0.02, 0.982$	$0.53 \pm 0.05, 0.625$
EIF	$0.74 \pm 0.15, 0.93$	$0.57 \pm 0.11, 0.67$
QIF	$0.90 \pm 0.05, 0.985$	$0.55 \pm 0.05, 0.625$

Table 6. Table of standard deviations on the N-MNIST dataset. Each cell reports the standard deviation calculated across all the experiments per each neuron and each task. Values are rounded to the closest third decimal. For each task, we highlighted in bold the highest values.

Neuron model	0 vs 1	2 vs 9	3 vs 7	4 vs 8
LIF	**0.092**	0.063	**0.068**	**0.061**
EIF	0.054	**0.074**	0.056	0.024
QIF	0.076	0.069	0.065	0.021

the N-MNIST, overall, the accuracy levels drastically increase. The LIF model accuracy grows by nearly 20 percentage points on average; however, the most striking increase is the accuracy of the QIF model, which gains 33 percentage points on average and 37.5 in the best case. This not only highlights the sensitivity of neuron models to their hyper-parameters, but also confirms that neurons with more complex dynamics can perform just as well as simpler ones.

In the case of the 'HC vs RHW' task, instead, we see a slightly different trend. In the first place, the results are surprisingly worse than those obtained through hand-picked parameters. We hypothesize that this is because the optimization system required more iterations to find a good set of hyper-parameters. As stated above, we allowed the same number of optimization iterations as in the case of the N-MNIST dataset. However, given the higher complexity of the features present in the DVS Gestures dataset, it might have been necessary to allow more. Secondly, although still struggling to achieve higher accuracy levels, the SNNs employing QIF and EIF averagely outperform those with the LIF. This confirms the results obtained using the same hyper-parameters and strengthen the hypothesis that richer dynamics can be beneficial when employed on data with a richer set of temporal features.

Another point worth considering is the variability of the results obtained. Spanning from the N-MNIST-based tasks to the gestures-based ones, the different neuron models demonstrate accuracy levels with a standard deviation of up to 15 percentage points. We hypothesize that this effect is caused by the order in which data is presented to the system in relation to the STDP learning rule which, as reported at the beginning of this section, is sensitive to such order. We also observe that the EIF model has higher fluctuations on average, which could possibly reflect a higher sensitivity to this effect.
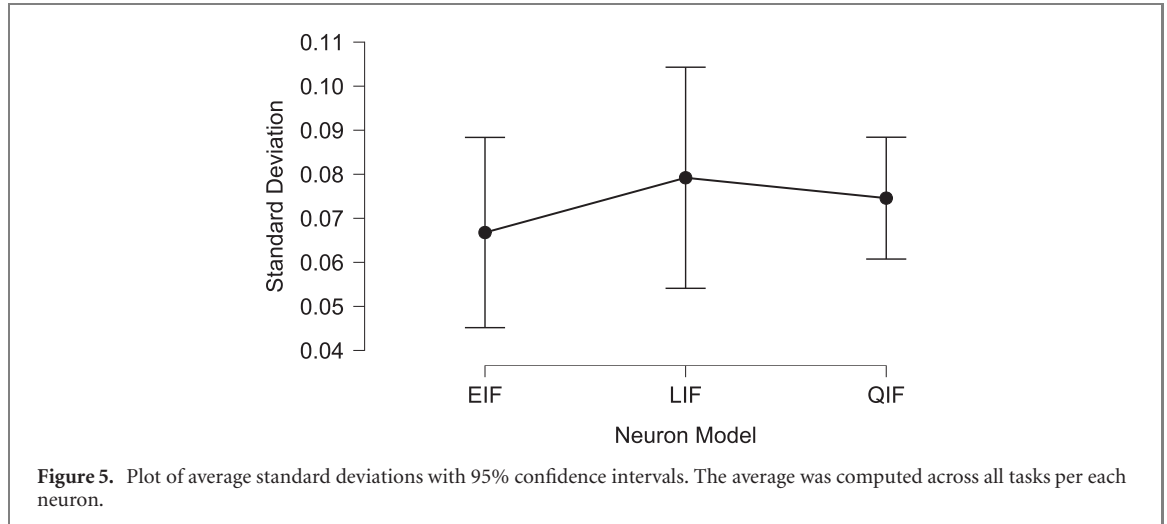
4.3. Sensitivity to data presentation order

In the previous paragraph, we hinted at the possibility of EIF neuron model being particularly sensitive to the presentation order of the data as a result of being trained through STDP. Driven by this, in this section we perform a series of controlled experiments aimed at studying this possibility.

In order to assess whether any neuron is more sensitive to the order of the data, we must constrain some parameters of the experiment to ensure the final results do not depend on these. Since we use a homogeneous set of neurons (they all share the same parameterization), the only varying elements are the convolutional weights (randomly initialized), and the order of the data itself. Hence, consistently with our previous methodology, we design a set of 11 experiments per task and neuron model in which the initialization of weights was fixed using a seed for the random number generator. Further to this, we allow each neuron model to process data presented in the same order. In other words, we evaluate each neuron model using the same order of data before changing to an other order for a total of 11 times per task. By doing so, we rule out the possibility of any neuron model displaying certain sensitivity levels as a result of fortunate/unfortunate randomization of the data. As a measure of the network sensitivity to the presentation order of the data we utilize the standard deviation of the accuracy across all the experiments per each neuron and task. These are reported in tables 6 and 7. We report the average sensitivity with 95% confidence intervals for each neuron model in figure 5.

Table 7. Table of standard deviations on the DVS gestures dataset. Each cell reports the standard deviation calculated across all the experiments per each neuron and each task. Values are rounded to the closest third decimal. For each task, we highlighted in bold the highest values.

Neuron model	HC vs RHW	RACW vs AG	RHCW vs AR	LACW vs AD
LIF	0.086	0.051	**0.122**	0.028
EIF	0.049	0.049	0.114	0.032
QIF	**0.088**	**0.079**	0.101	**0.074**

Figure 5. Plot of average standard deviations with 95% confidence intervals. The average was computed across all tasks per each neuron.

We perform a one-way analysis of variance (ANOVA) to verify whether any neuron is significantly more sensitive than others. The test assumptions were checked. Levene's test was non-significant ($p = 0.256$), indicating that the assumption of homogeneity of variance was not violated. Normality was checked with a Q–Q plot. No deviations were noted. We found no significant difference among the three neuron models in sensitivity to the order of the data, $F(2, 21) = 0.514$, $p = 0.605$, $\eta_p^2 = 0.047$. In the context of our experiments, these findings indicate that using either neuron model does not increase nor decrease the sensitivity to the presentation order of the data. Instead, such sensitivity is probably only due to the use of STDP as a learning rule.

5. Discussion

5.1. Implications of using different neuron models

The usage of spiking neuron models has some inherent implications on the machine learning pipeline from the implementation and the theoretical points of view.

Concerning the implementation, spiking neuron come with a whole set of hyper-parameters to tune. Considering the LIF, the simplest version requires a single parameter (the time constant or a leakage term), but other implementations might include up to five different parameters, such as the refractory period or the time-step size. If we switch to the QIF or the EIF, there are at least two new and non-optional parameters to consider (see 2.1).

Determining a good set of hyper-parameters is a non-trivial task [75]. Although in the NM field a lot of inspiration is taken from the human brain, it is not possible to simply assume that the same parameters that work in such a complex system would still be applicable in a simplification such as an SNN. We hence need to tweak parameters for our need or, alternatively, to define a parameter optimization strategy that does that heuristically in an automated way. However, also the latter solution often requires to make guesses about the domain in which parameters can vary and it requires a long time to compute. Furthermore, hyper-parameters can be correlated in some way, thus making both the hand-tuning and the automated optimization process more difficult. As a result, from an implementation point of view, using neuron models that require more hyper-parameters can significantly increase the usage complexity.

From a theoretical point of view, using different neurons or varying the parameters means to open to different non-linear dynamics and excitability patterns. Figures 6 and 7 provide a visual understanding of these
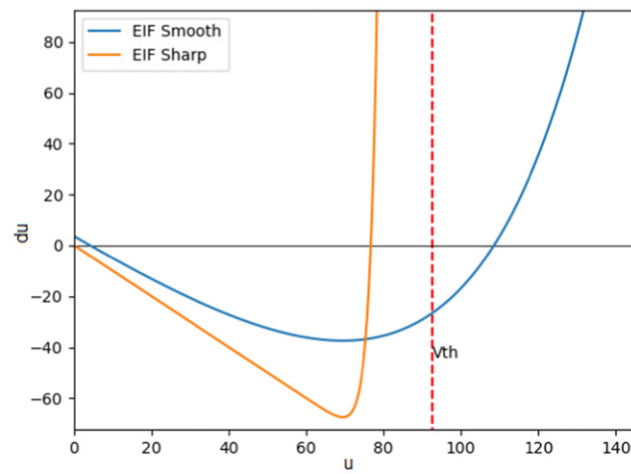
Figure 6. Example of membrane potential dynamics of the spiking neuron models. In both figures, the *y*-axis represents the variation of the membrane potential d*u*, while the *x*-axis represents the value of the membrane potential itself *u*. Figure 6(a) compares the three spiking neurons, whereas figure 6(b) is an example of how varying the sharpness parameter Δ_T can affect the dynamics of the EIF.

differences. In the LIF, the membrane potential updates depend linearly on the previous state of the membrane potential itself. The EIF manifests a similar relationship up to certain values of membrane potential (Θ_{rh}), after which the relationship assumes a more non-linear (exponential) aspect. The QIF loses any linear relationship in favor of a quadratic one. These dynamics play a role on the excitability (regions) of a neuron [2]. For instance, an EIF with a smooth exponential term (blue line in figure 6(b)) will receive more mitigated updates throughout (slow-forgetting neuron), whereas an EIF with a sharp exponential term (orange line) will receive more negative updates up to the cutoff threshold Θ_{rh} and then highly positive ($+\infty$) ones, thus immediately reaching the firing threshold V_{th}. Hence, the second example would be a fast-forgetting neuron, but the cut-off threshold will act as an early firing threshold, as any subsequent update would bring the membrane potential up and above the actual firing threshold. A similar example is reported in figure 7, where a change in the time-constant τ_m makes a LIF neuron forget faster or slower. This in turn has effects on the excitability of the neuron and its firing pattern.

The different firing abilities discussed above need to be considered within the context where the neurons are placed. If we consider the case of a homogeneous SCNN that is trained on a dataset in which the temporal distribution of events is similar for every sample, it might be pointless to consider having a wide range of excitability patters as more complex neurons have. Indeed, the increased amount of parameters would make it more difficult to find the right excitability that works well with that data. At the same time, they would likely come at a higher computational and power cost. Conversely, if the dataset is considerably diverse in terms of temporal distribution of events, it would arguably prove useful to have a broad range of excitability patterns
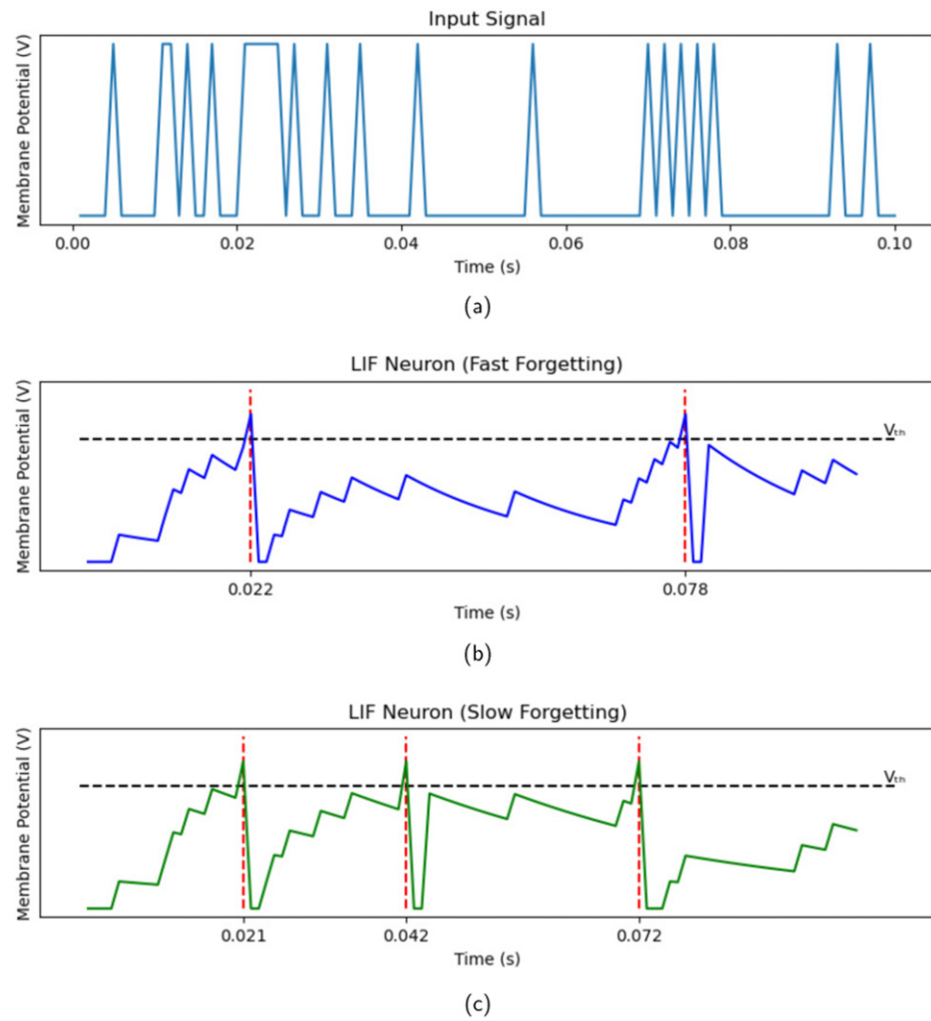
Figure 7. Example of changing hyper-parameters in a LIF neuron. The 'fast forgetting' neuron (smaller time constant τ_m) (figure 7(b)) can only spike twice in response to the input (figure 7(a)). The 'slow forgetting' one (figure 7(c)) can fire three times and maintains a higher membrane potential throughout. Note that the 'slow forgetting' neuron also fires earlier, i.e. it requires less (close) spikes to reach the threshold. Both the neurons have a refractory period of 2 ms.

to choose from. Therefore, an heterogeneous network of spiking neurons would likely be able to learn better or simply more features. In this context, employing more complex neurons with variable parameters can be significantly beneficial.

5.2. Temporal features and neuron performance

In our experiments, we used a extremely simple homogeneous SCNN to perform a simple classification task on a simple subset of the N-MNIST and DVS gesture datasets. The N-MNIST dataset, although natively event-based, is not a naturally dynamic dataset. The original data, the MNIST handwritten digits, are static images that do not contain temporal dynamical features. As such, the temporal features that are instead present in the N-MNIST are crafted. Furthermore, these dynamics are obtained by moving a DVS camera using the same sequence of movements with the same timing. Figure 2(c) shows that, as a result, the distribution of events throughout each sample present in the dataset is roughly the same. This means that the temporal features are not different from one another and are hence not discriminative of different classes of samples. Indeed, as discussed in the previous paragraph, using a homogeneous SCNN was enough to achieve reasonably high accuracy levels, despite the lack of diversity in the dynamics of the embedded neurons.

Concerning the performance in such a homogeneous settings, in our consideration of three single-variable neuron models we found that all of them have the ability to perform well. The difference however is in the cost of using one neuron rather than the other. From our experiments, we found that when hand-tuning parameters, the LIF neuron achieved averagely high accuracy levels, with EIF neuron being better at times. Using a set of optimized hyper-parameters, we observed a considerable improvement in the overall classification accuracy with the QIF achieving a 98.5% accuracy in the best case, thus surpassing its counterparts. The same QIF model performed rather poorly when using non-optimized parameters. As mentioned in section 4.2, this highlights

the fact that, despite the data displaying simple spatio-temporal features, more complex neurons are still able to perform well. The cost for achieving such results can, however, become rather high.

When employing gestures data, the situations is slightly different. In this case, as depicted in figure 3, there is no recurring distribution of events across different classes. Instead, events are distributed throughout the whole time domain. Each class of gestures has a distribution of events that varies with respect to the others, even more because of the fact that different actions require a different time to be executed. Thus, the temporal features in this dataset are more important and diverse. As a matter of fact, this is also shown by the results obtained using the same network as in the previous case. Here, although the performance gain is still modest, the EIF and QIF neurons steadily attain better classification accuracies than the LIF model. Since we used the same setting for all the experiments, this is likely traceable to the aforementioned differences in temporal features, which are now more diverse and complex than those in the N-MNIST dataset.

5.3. Temporal features and depth of the network

The matter of temporal variety in the features being better represented by more complex dynamics opens up further questions as to their use in SNNs. Indeed, if we consider a hierarchical NN, the deepest layers normally learn more abstract representations, whereas the early layers typically learn to distinguish simple patterns, such as edges or corners [44]. This is easily conceivable when thinking about spatial features. For example, when a set of lines is recognized in the early layers of a CNN, these could later be understood to be a square, and further down the network as a house. Although it can be more difficult to imagine, when we include the time dimension, similar scenarios can arise where features relate temporally other than spatially. It thus seems straightforward that the temporal relationships might vary in complexity in different stages of the network depending on the task at hand.

We showed that the use of more complex neuron models improves the performance on more complex tasks at the level of one layer. When combining several of these layers in a hierarchical network more uses of their non-linear dynamics could arise, as they would combine several spatio-temporal features built up in previous stages to understand compound featural patterns.

6. Conclusion

In this work, we have considered a simple unsupervised SCNN and analyzed the effect on the overall performance of changing the underlying neuron models. Because they were not previously present in the SpykeTorch framework, we implement the spiking neurons and make the code available. We firstly draw a set of four binary classification tasks using 4 couples of classes from the N-MNIST dataset on which we repeatedly train and evaluate our SCNN. Experimental results on these tasks show that all three neuron models can achieve top-level accuracies, albeit the more complex ones require more fine-tuning. In a second instance, we consider the DVS gestures dataset, which exposes a richer set of features from both the visual and temporal points of view. In this case, the EIF and QIF steadily outperform the LIF on all four tasks drawn from this dataset. Further to this, we analyzed the sensitivity of the SCNN using each neuron model to the order of presentation of the data. Our analysis shows that none of them implies a statistically relevant difference in terms of sensitivity, and that such sensitivity is probably only relatable to the use of STDP, in our experiments. Collectively, our results show that accurately selecting the neuron model employed in an NM pipeline improves its performance, and that this selection should be driven by considering the complexity of the spatio-temporal features that the layer in the network will have to understand. Furthermore, it highlights that further research aimed at unveiling the role of the dynamics of neuron models in deep hierarchical learning would be highly beneficial to close the gap between conventional DL approaches and SNNs. Other future studies could consist of analysing further relationships between the neuron models and other components of the learning pipeline, such as the neural network architecture, and the learning rule. Furthermore, it would be interesting to investigate the result of using different spiking neurons on the synaptic efficacies when using STDP as a learning rule.

Acknowledgments

Data availability statement

No new data were created or analysed in this study.

ORCID iDs

Davide L Manna ⬤ https://orcid.org/0000-0001-8963-5050
Alex Vicente-Sola ⬤ https://orcid.org/0000-0002-2370-6562
Paul Kirkland ⬤ https://orcid.org/0000-0001-5905-6816
Trevor J Bihl ⬤ https://orcid.org/0000-0003-2431-2749
Gaetano Di Caterina ⬤ https://orcid.org/0000-0002-7256-0897

References

[1] Schuman C D, Potok T E, Patton R M, Birdwell J D, Dean M E, Rose G S and Plank J S 2017 A survey of neuromorphic computing and neural networks in hardware (arXiv: 1705.06963)

[2] Gerstner W, Kistler W M, Naud R and Paninski L 2009 *Neuronal Dynamics* (Cambridge: Cambridge University Press)

[3] Schemmel J, Billaudelle S, Dauer P and Weis J 2020 Accelerated analog neuromorphic computing (arXiv:abs/2003.11996)

[4] Benjamin B V *et al* 2014 Neurogrid: a mixed-analog-digital multichip system for large-scale neural Simulations *Proc. IEEE* **102** 699–716

[5] Davies M *et al* 2018 Loihi: a neuromorphic Manycore processor with on-chip learning *IEEE Micro* **38** 82–99

[6] Merolla P, Arthur J, Akopyan F, Imam N, Manohar R and Modha D S 2011 A digital neurosynaptic core using embedded crossbar memory with 45pj per spike in 45nm *2011 IEEE Custom Integrated Circuits Conf. (CICC)* pp 1–4

[7] Cassidy A S *et al* 2013 Cognitive computing building block: a versatile and efficient digital neuron model for neurosynaptic cores *2013 Int. Joint Conf. Neural Networks (IJCNN)* pp 1–10

[8] Akopyan F *et al* 2015 TrueNorth: design and tool flow of a 65 mW 1 million neuron programmable neurosynaptic chip *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.* **34** 1537–57

[9] Yang S, Wang J, Hao X, Li H, Wei X, Deng B and Loparo K A 2022 BiCoSS: toward large-scale cognition brain with multigranular neuromorphic architecture *IEEE Trans. Neural Netw. Learning Syst.* **33** 2801–15

[10] Furber S B, Galluppi F, Temple S and Plana L A 2014 The spinnaker project *Proc. IEEE* **102** 652–65

[11] Orchard G, Frady E P, Rubin D B D, Sanborn S, Bam Shrestha S, Sommer F T and Davies M 2021 Efficient neuromorphic signal processing with Loihi 2 *2021 IEEE Workshop on Signal Processing Systems (SiPS)* (Piscataway, NJ: IEEE) pp 254–9

[12] Vicente-Sola A, Manna D L, Kirkland P, Di Caterina G and Bihl T 2022 Keys to accurate feature extraction using residual spiking neural networks *Neuromorphic Computing and Engineering* **2** 044001

[13] Aamir S, Stradmann Y, Muller P, Pehle C, Hartel A, Grubl A, Schemmel J and Meier K 2018 An accelerated LIF neuronal network array for a large-scale mixed-signal neuromorphic architecture *IEEE Trans. Circuits Syst.* I **65** 4299–312

[14] Diamond A, Nowotny T and Schmuker M 2016 Comparing neuromorphic solutions in action: implementing a bio-inspired solution to a benchmark classification task on three parallel-computing platforms *Front. Neurosci.* **9** 491

[15] Friedl K E, Voelker A R, Peer and Eliasmith C 2016 Human-inspired neurorobotic system for classifying surface textures by touch *IEEE Robot. Autom. Lett.* **1** 516–23

[16] Hunsberger E and Eliasmith C 2015 Spiking Deep Networks with Lif Neurons (arXiv: 1510.08829)

[17] Göltz J *et al* 2021 Fast and energy-efficient neuromorphic deep learning with first-spike times *Nat. Mach. Intell.* **3** 823–35

[18] Mozafari M, Kheradpisheh S R, Masquelier T, Nowzari-Dalini A and Ganjtabesh M 2018 First-spike-based visual categorization using reward-modulated STDP *IEEE Trans. Neural Netw. Learning Syst.* **29** 6178–90

[19] Stromatias E, Neil D, Galluppi F, Pfeiffer M, Liu S-C and Furber S 2015 Scalable energy-efficient, low-latency implementations of trained spiking deep belief networks on SpiNNaker *2015 Int. Joint Conf. Neural Networks (IJCNN)* (Piscataway, NJ: IEEE)

[20] Mostafa H 2018 Supervised learning based on temporal coding in spiking neural networks *IEEE Trans. Neural Netw. Learning Syst.* **29** 3227–35

[21] Chaturvedi S, Khurshid A A and Boudjelal M 2012 Image segmentation using leaky integrate-and-fire model of spiking neural network *Int. J. Wisdom Based Comput.* **2** 21–8

[22] Jiang C, Yang L and Zhang Y 2021 A spiking neural network with spike-timing-dependent plasticity for surface roughness analysis *IEEE Sens. J.* **22** 438–45

[23] Yang S, Wang J, Zhang N, Deng B, Pang Y and Azghadi M R 2022 CerebelluMorphic: large-scale neuromorphic model and architecture for supervised motor learning *IEEE Trans. Neural Netw. Learning Syst.* **33** 4398–412

[24] Patiño-Saucedo A, Rostro-Gonzalez H, Serrano-Gotarredona T and Linares-Barranco B 2020 Event-driven implementation of deep spiking convolutional neural networks for supervised classification using the spinnaker neuromorphic platform *Neural Netw.* **121** 319–28

[25] Fardet T, Ballandras M, Bottani S, Métens S and Monceau P 2018 Understanding the generation of network bursts by adaptive oscillatory neurons *Front. Neurosci.* **12** 41

[26] Taherkhani K and Shoorehdeli M A 2017 An artificial neural network based on izhikevich neuron model *2017 Iranian Conf. Electrical Engineering (ICEE)* (Piscataway, NJ: IEEE)

[27] Chaturvedi S, Titre R N and Sondhiya N 2014 Review of handwritten pattern recognition of digits and special characters using feed forward neural network and izhikevich neural model *2014 Int. Conf. Electronic Systems, Signal Processing and Computing Technologies* (Piscataway, NJ: IEEE)

[28] Vazquez R A 2011 Training spiking neural models using cuckoo search algorithm *2011 IEEE Congress of Evolutionary Computation (CEC)* (Piscataway, NJ: IEEE)

[29] Markram H, Toledo-Rodriguez M, Wang Y, Gupta A, Silberberg G and Wu C 2004 Interneurons of the neocortical inhibitory system *Nat. Rev. Neurosci.* **5** 793–807

[30] Connors B W and Gutnick M J 1990 Intrinsic firing patterns of diverse neocortical neurons *Trends Neurosci.* **13** 99–104

[31] Dauth S, Maoz B M, Sheehy S P, Hemphill M A, Murty T, Macedonia M K, Greer A M, Budnik B and Parker K 2017 Neurons derived from different brain regions are inherently different *in vitro*: a novel multiregional brain-on-a-chip *J. Neurophysiol.* **117** 1320–41 PMID: 28031399

[32] Orchard G, Jayawant A, Cohen G K and Thakor N 2015 Converting static image datasets to spiking neuromorphic datasets using saccades *Front. Neurosci.* **9** 437

[33] Amir A *et al* 2017 A low power, fully event-based gesture recognition system *2017 IEEE Conf. Computer Vision and Pattern Recognition (CVPR)* pp 7388–97

[34] Mozafari M, Ganjtabesh M, Nowzari-Dalini A and Masquelier T 2019 Spyketorch: efficient simulation of convolutional spiking neural networks with at most one spike per neuron *Front. Neurosci.* **13** 625

[35] Lapicque L 1907 Recherches quantitatives sur l'excitation electrique des nerfs traitée comme une polarization *J. Phys. Path. General* **9** 620–35

[36] Hodgkin A L and Huxley A F 1952 A quantitative description of membrane current and its application to conduction and excitation in nerve *J. Physiol.* **117** 500–44

[37] Gerstner W and Kistler W M 2002 *Spiking Neuron Models* (Cambridge: Cambridge University Press)

[38] Gerstner W, Ritz R and van Hemmen J 1993 A biologically motivated and analytically soluble model of collective oscillations in the cortex *Biol. Cybern.* **68** 363–74

[39] van Vreeswijk C and Sompolinsky H 1996 Chaos in neuronal networks with balanced excitatory and inhibitory activity *Science* **274** 1724–6

[40] Kawato M and Gomi H 1992 A computational model of four regions of the cerebellum based on feedback-error learning *Biol. Cybern.* **68** 95–103

[41] Zipser D, Kehoe B, Littlewort G and Fuster J 1993 A spiking network model of short-term active memory *J. Neurosci.* **13** 3406–20

[42] Maass W 1996 Lower bounds for the computational power of networks of spiking neurons *Neural Comput.* **8** 1–40

[43] Maass W 1997 Networks of spiking neurons: the third generation of neural network models *Neural Netw.* **10** 1659–71

[44] Goodfellow I, Bengio Y and Courville A 2016 *Deep Learning* (Cambridge, MA: MIT Press)

[45] Sharma S, Sharma S and Athaiya A 2017 Activation functions in neural networks *Towards data science* **6** 310–6

[46] Maas A L *et al* 2013 Rectified nonlinearities improve neural network acoustic model *Int. Conf. on Machine Learning (ICML)* vol 30

[47] He K, Zhang X, Ren S and Sun J 2015 Delving deep into rectifiers: surpassing human-level performance on imagenet classification *2015 IEEE Int. Conf. Computer Vision (ICCV)* pp 1026–34

[48] Clevert D-A, Unterthiner T and Hochreiter S 2015 Fast and accurate deep network learning by exponential linear units (elus) (arXiv:1511.07289)

[49] Ramachandran P, Zoph B and Le Q V 2018 Searching for activation functions (arXiv:abs/1710.05941)

[50] Pedamonti D 2018 Comparison of non-linear activation functions for deep neural networks on mnist classification task (arXiv:1804.02763)

[51] Eger S, Youssef P and Gurevych I 2019 Is it time to swish? comparing deep learning activation functions across NLP tasks (arXiv:1901.02671)

[52] Ding B, Qian H and Zhou J 2018 Activation functions and their characteristics in deep neural networks *2018 Chinese Control and Decision Conf. (CCDC)* pp 1836–41

[53] Zheng Q, Yang M, Tian X, Wang X and Wang D 2020 Rethinking the role of activation functions in deep convolutional neural networks for image classification *Eng. Lett.* **28** 80–92

[54] Goyal M, Goyal R, Venkatappa Reddy P and Lall B 2020 *Activation Functions* (Berlin: Springer) pp 1–30

[55] Jagtap A D, Kawaguchi K and Karniadakis G E 2020 Adaptive activation functions accelerate convergence in deep and physics-informed neural networks *J. Comput. Phys.* **404** 109136

[56] Sboev A, Rybka R, Serenko A, Vlasov D, Kudryashov N and Demin V 2018 To the role of the choice of the neuron model in spiking network learning on base of spike-timing-dependent plasticity *Proc. Comput. Sci.* **123** 432–9

[57] Izhikevich E M 2004 Which model to use for cortical spiking neurons? *IEEE Trans. Neural Netw.* **15** 1063–70

[58] Long L and Fang G 2010 A review of biologically plausible neuron models for spiking neural networks *AIAA Infotech@Aerospace* (American Institute of Aeronautics and Astronautics)

[59] Barton A, Volna E and Kotyrba M 2018 The application perspective of Izhikevich spiking neural model—the initial experimental study *Recent Advances in Soft Computing* (Berlin: Springer) pp 223–32

[60] Kumar G, Aggarwal V, Thakor N v, Schieber M and Kothare M V 2010 Optimal parameter estimation of the izhikevich single neuron model using experimental inter-spike interval (isi) data 3586–91

[61] Teeter C *et al* 2018 Generalized leaky integrate-and-fire models classify multiple neuron types *Nat. Commun.* **9** 709

[62] Jolivet R, Rauch A, Lüscher H-R and Gerstner W 2005 Integrate-and-fire models with adaptation are good enough *Advances in Neural Information Processing Systems* vol 18 ed Y Weiss, B Schölkopf and J Platt (Cambridge, MA: MIT Press)

[63] Grzyb B J, Chinellato E, Wojcik G M and Kaminski W A 2009 Which model to use for the liquid state machine? *2009 Int. Joint Conf. Neural Networks* (Piscataway, NJ: IEEE)

[64] Traub R D, Wong R K, Miles R and Michelson H 1991 A model of a Ca$_3$ hippocampal neuron incorporating voltage-clamp data on intrinsic conductances *J. Neurophysiol.* **66** 635–50

[65] Halnes G, Augustinaite S, Heggelund P, Einevoll G T and Migliore M 2011 A multi-compartment model for interneurons in the dorsal lateral geniculate nucleus *PLoS Comput. Biol.* **7** e1002160

[66] Shimada N and Torikai H 2015 A novel asynchronous cellular automaton multicompartment neuron model *IEEE Trans. Circuits Syst. II* **62** 776–80

[67] Yang S, Gao T, Wang J, Deng B, Azghadi M R, Lei T and Linares-Barranco B 2022 SAM: a unified self-adaptive multicompartmental spiking neuron model for learning with working memory *Front. Neurosci.* **16** 850945

[68] Yang S, Linares-Barranco B and Chen B 2022 Heterogeneous ensemble-based spike-driven few-shot online learning *Front. Neurosci.* **16** 850932

[69] Izhikevich E M 2010 *Dynamical Systems in Neuroscience* (Cambridge, MA: MIT Press)

[70] FitzHugh R 1961 Impulses and physiological states in theoretical models of nerve membrane *Biophys. J.* **1** 445–66

[71] Nagumo J, Arimoto S and Yoshizawa S 1962 An active pulse transmission line simulating nerve axon *Proc. IRE* **50** 2061–70

[72] Hindmarsh J L and Rose R M 1984 A model of neuronal bursting using three coupled first order differential equations *Proc. R. Soc. London* B **221** 87–102

[73] Morris C and Lecar H 1981 Voltage oscillations in the barnacle giant muscle fiber *Biophys. J.* **35** 193–213

[74] Burkitt A N 2006 A review of the integrate-and-fire neuron model: I. homogeneous synaptic input *Biol. Cybern.* **95** 1–19

[75] Gerstner W and Naud R 2009 How good are neuron models? *Science* **326** 379−80

[76] Fourcaud-Trocmé N, Hansel D, van Vreeswijk C and Brunel N 2003 How spike generation mechanisms determine the neuronal response to fluctuating inputs *J. Neurosci.* **23** 11628−40

[77] Ermentrout G B and Kopell N 1986 Parabolic bursting in an excitable system coupled with a slow oscillation *SIAM J. Appl. Math.* **46** 233−53

[78] Brette R and Gerstner W 2005 Adaptive exponential integrate-and-fire model as an effective description of neuronal activity *J. Neurophysiol.* **94** 3637−42

[79] Izhikevich E M 2003 Simple model of spiking neurons *IEEE Trans. Neural Netw.* **14** 1569−72

[80] Cheng X, Hao Y, Xu J and Xu B 2020 Lisnn: improving spiking neural networks with lateral interactions for robust object recognition *Proc. 29th Int. Joint Conf. Artificial Intelligence, IJCAI-20. Int. Joint Conf. Artificial Intelligence Organization* vol 7 ed C Bessiere pp 1519−25

[81] Manna D L, Vicente-Sola A, Kirkland P, Bihl T J and Di Caterina G 2022 Frameworks for SNNs: A Review of Data Science-Oriented Software and an Expansion of Spyketorch [Manuscript submitted for publication]

[82] Masquelier T and Thorpe S J 2007 Unsupervised learning of visual features through spike timing dependent plasticity *PLoS Comput. Biol.* **3** e31

[83] Kirkland P, Di Caterina G, Soraghan J and Matich G 2020 Spikeseg: spiking segmentation via stdp saliency mapping *2020 Int. Joint Conf. Neural Networks (IJCNN)* pp 1−8

[84] Kirkland P, Manna D L, Vicente-Sola A and Di Caterina G 2021 Unsupervised spiking instance segmentation on event data using STDP (arXiv:abs/2111.05283)

[85] Taherkhani A, Belatreche A, Li Y, Cosma G, Maguire L P and McGinnity T M 2020 A review of learning in biologically plausible spiking neural networks *Neural Netw.* **122** 253−72

[86] Diehl P and Cook M 2015 Unsupervised learning of digit recognition using spike-timing-dependent plasticity *Front. Comput. Neurosci.* **9** 99

[87] Iyer L R and Basu A 2017 Unsupervised learning of event-based image recordings using spike-timing-dependent plasticity *2017 Int. Joint Conf. Neural Networks (IJCNN)* pp 1840−6

[88] Falkner S, Klein A and Hutter F 2018 BOHB: robust and efficient hyperparameter optimization at scale *Int. Conf. on Machine Learning (ICML)* (Proc. Machine Learning Research) (PMLR) vol 80 pp 1437−1446