Long Short-Term Memory Spiking Networks and Their Applications

Ali Lotfi Rezaabad

alotfi@utexas.edu

The University of Texas at Austin
Austin, Texas

Sriram Vishwanath

sriram@austin.utexas.edu

The University of Texas at Austin
Austin, Texas

ABSTRACT

Recent advances in event-based neuromorphic systems have resulted in significant interest in the use and development of spiking neural networks (SNNs). However, the non-differentiable nature of spiking neurons makes SNNs incompatible with conventional backpropagation techniques. In spite of the significant progress made in training conventional deep neural networks (DNNs), training methods for SNNs still remain relatively poorly understood. In this paper, we present a novel framework for training recurrent SNNs. Analogous to the benefits presented by recurrent neural networks (RNNs) in learning time series models within DNNs, we develop SNNs based on long short-term memory (LSTM) networks. We show that LSTM spiking networks learn the timing of the spikes and temporal dependencies. We also develop a methodology for error backpropagation within LSTM-based SNNs. The developed architecture and method for backpropagation within LSTM-based SNNs enable them to learn long-term dependencies with comparable results to conventional LSTMs. Code is available on github; <https://github.com/AliLotfi92/SNNLSTM>

ACM Reference Format:

Ali Lotfi Rezaabad and Sriram Vishwanath. 2020. Long Short-Term Memory Spiking Networks and Their Applications. In *International Conference on Neuromorphic Systems 2020 (ICONS 2020)*, July 28–30, 2020, Oak Ridge, TN, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3407197.3407211>

1 INTRODUCTION

The development and successful training of deep neural networks (DNNs) has resulted in breakthrough results in different application areas such as computer vision and machine learning [19, 21, 33]. Although neural networks are inspired by neurons in the nervous system, it is known that learning and computation in nervous system is mainly based on event-based spiking computational units [7]. Accordingly, spiking neural networks (SNNs) have been proposed to better mimic the capabilities of biological neural networks. Although SNNs can represent the underlying spatio-temporal behavior of biological neural networks, they received much less attention, due to difficulties in training since spikes in general are not

differentiable and gradient-based methods cannot be used directly for training.

SNNs, similar to DNNs are formed of multiple layers and several neurons per layer. They differ in functionality, however, with SNNs sharing *spikes* rather than floating point values.

In general, DNNs and SNNs can be reduced to optimized ASICs and/or parallelized using GPUs. Due to temporal sparsity, ASIC implementations of SNNs are found to be far more energy and resource efficient, with neuromorphic chips emerging that possess high energy efficiency, including Loihi [6], SpiNNaker [10] and others [27, 28]. This energy efficiency, along with their relative simplicity in inference make SNNs attractive, so long as they can be trained efficiently, and perform in a manner similar to DNNs.

Through this paper, we focus on recurrent SNNs. Similar to recurrent DNNs, recurrent SNNs are a special class of SNNs that are equipped with an internal memory which is managed by the network itself. This additional storage gives them the power to process sequential dataset. Hence, they are popular for different tasks including speech recognition and language modeling.

Despite the substantial literature on training SNNs, the domain, especially recurrent SNNs, is still in its infancy when compared to our understanding of training mechanisms for DNNs. A significant portion of SNN-training literature has focused on training feed-forward SNNs with one layer networks [14, 24]. Recently, some developments enabled training multi layer SNNs [31], nonetheless, training recurrent SNNs is still in an incipient stage.

Recently, [31] utilized spike responses based on kernel functions for every neuron to capture the temporal dependencies of spike trains. Although this method successfully captures the temporal dependency between spikes, kernel-based computations are costly. Moreover, the need for convolution operation over time makes them inefficient to be applied to recurrent SNNs.

Our contributions. We present a new framework for designing and training recurrent SNNs based on long short-term memory (LSTM) units. Each LSTM unit includes three different gates: forget gate that helps to dismiss useless information, input gate monitors the information entering the unit, and output gate that forms the outcome of the unit. Indeed, LSTM [15] and its variants [13] are special cases of recurrent neural networks (RNNs) that, in part, help address the vanishing gradient problem. LSTMs are considered particularly well-suited for time series and sequential datasets. In this paper, we leverage this capability within SNNs to propose LSTM-based SNNs that are capable of sequential learning. We propose a novel backpropagation mechanism and architecture in this paper which make it possible to achieve better performance than existing recurrent SNNs that is comparable with conventional LSTMs. In addition, our approach does not require a convolutional mechanism

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICONS 2020, July 28–30, 2020, Oak Ridge, TN, USA

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-8851-1/20/07...\$15.00

<https://doi.org/10.1145/3407197.3407211>

over time, resulting in a lower-complexity training mechanism for recurrent SNNs compared to the feedforward neural network kernel-based approaches.

We study the performance and dynamics of our proposed architecture through empirical evaluations on various datasets. First, we start with a toy datasets, and then follow by benchmark language modeling and speech recognition datasets which provide more structured temporal dependencies. Additionally, our approach achieves better test accuracy compared to the existing literature using a simple model and network. Further, we also show that such an LSTM SNN performs well on the larger and more complex sequential EMNIST dataset [4]. Finally, we evaluate the capability of the proposed recurrent SNNs in natural-language generation which reveals one of many interesting applications of SNNs.

2 RELATED WORK

In general, existing approaches for training SNNs can be subdivided into *indirect training* and *direct training* categories. *Indirect training* of SNNs refers to those approaches that train a conventional DNN using existing approaches and then associate/map the trained output to the desired SNN. Such a mechanism can be fairly general and powerful, but it can be limiting as the SNN obtained depends heavily on the associated DNN. In particular, [9] presents a framework where they optimize the probability of spiking on a DNN, and then transfer the optimized parameters into the SNN. Further literature has been developed on this framework by adding noise to the associated activation function [22], constraining the synaptics' strengths (the network's weights and biases) [8], and utilizing alternate transfer functions [26].

To enable *direct training* of SNNs, SpikeProp [2] presents a pioneering supervised temporal learning algorithm. Here, the authors simulate the dynamics of neurons by leveraging an associated spike response model (SRM) [11]. In particular, SpikeProp and its associated extensions [3, 29] update the weights in accordance with the actual and target spiking times using gradient descent. However, the approach is challenging to be applied to benchmark tasks. To partially address this, improvements on SpikeProp have been developed, including MuSpiNN [12], and Resilient propagation [23]. More recently, [17] presents a two-level backpropagation algorithm for training SNNs, and [31] presents a framework for training SNNs where both weights and delays are optimized simultaneously. Additionally, these frameworks apply a kernel function for every neuron, which might be a memory-intensive and time-consuming operation, especially for recurrent SNNs.

Perhaps the most related to our work is the recent work in [1]. Similarly, the authors propose using LSTM units and in relation with the algorithm in [16] to assure that the neurons in LSTM units output either 1 or -1. For training, they approximate the gradient of the spike activation with the piecewise linear function $\max\{0, 1 - |u|\}$, where u is the output of the neuron before the activation (so-called neurons' membrane potential). In this paper, however, we relaxed the gradient of the spike activation with a probability distribution. This relaxation provides more precise updates for the network at each iteration. Also authors in [5] have studied to remodel the architecture of LSTM to be admissible to cortical circuits which are similar to the circuits have been found in

Figure 1: An LSTM spiking unit composed of: 1- forget gate layer f_t , 2- input gate layer i_t , 3- output gate layer o_t , 4- modulated input g_t , 5- hidden state h_t , 6- unit state c_t .

nervous system. Indeed, they leverage the sigmoid function for all activations in LSTM. Further, [30] is an *indirect training* approach where they first run a conventional LSTM and then map it into spiking version.

There are bio-inspired approaches for training SNNs, including methods such as spike-time dependent plasticity (STDP) [32] for *direct training*. STDP is an unsupervised learning mechanism which mimics the human visual cortex. Although such biologically-inspired training mechanisms are of interest, they are also challenging to benchmark, and therefore, we focus on alternative *direct training* approaches in this paper.

3 OUR METHODOLOGY

3.1 LSTM Spiking Neural Networks

LSTM and its variants, a special class of RNNs, are popular due to their remarkable results in different sequential processing tasks, including long-range structures, i.e., natural language modeling and speech recognition. Indeed, LSTMs and in general RNNs are capable of capturing the temporal dependence of their input, while also addressing the vanishing gradient issue faced by other architectures.

Therefore, LSTM networks constitute a natural candidate to capture the temporal dependence a SNN models. The output value of a neuron before applying the activation is called its membrane potential, denoted as $u_n(t)$ for neuron n at time t , see Figure 2a.

We outline LSTM spiking unit's main elements in Figure 1. An LSTM spiking unit has three interacting gates and associated "spike" functions. Generally, spike activations $\sigma_1(u)$ and $\sigma_2(u)$ are applied to each of their associated neurons individually. These functions take neurons' membrane potential $u_n(t)$ and outputs either a spike or null at each time step.

Like conventional LSTMs, the core idea behind such an LSTM spiking unit is the unit state, c_t , which is a pipeline and manager of information flow between units. Indeed, this is done through collaborations of different gates and layers. Forget gate, denoted by f_t , decides what information should be dismissed. The input gate i_t , controls the information entering the unit, and another assisting layer on input, g_t , which is modulated by another spike activation $\sigma_2(u)$. Eventually, the output of the unit is formed based

on the output gate $\mathbf{o}_t$, and the unit state. More specifically, given a set of spiking inputs $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T\}$, the gates and states are characterized as follows:

$$\begin{aligned}
 \mathbf{f}_t &= \sigma_1(w_{f,h}\mathbf{h}_{t-1} + w_{f,x}\mathbf{x}_t + \mathbf{b}_{f,h} + \mathbf{b}_{f,x}), \\
 \mathbf{i}_t &= \sigma_1(w_{i,h}\mathbf{h}_{t-1} + w_{i,x}\mathbf{x}_t + \mathbf{b}_{i,h} + \mathbf{b}_{i,x}), \\
 \mathbf{g}_t &= \sigma_2(w_{g,h}\mathbf{h}_{t-1} + w_{g,x}\mathbf{x}_t + \mathbf{b}_{g,h} + \mathbf{b}_{g,x}), \\
 \mathbf{c}_t &= \mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \mathbf{g}_t, \\
 \mathbf{o}_t &= \sigma_1(w_{o,h}\mathbf{h}_{t-1} + w_{o,x}\mathbf{x}_t + \mathbf{b}_{o,h} + \mathbf{b}_{o,x}), \\
 \mathbf{h}_t &= \mathbf{o}_t \odot \mathbf{c}_t,
 \end{aligned} \tag{1}$$

where $\odot$ represents the Hadamard product, $\sigma_1(\cdot)$ and $\sigma_2(\cdot)$ are spike activations that map the membrane potential of a neuron, $u_n(t)$, to a spike if it exceeds the threshold value θ_1 and θ_2 , respectively. Throughout this paper, we assume two expressions: 1) wake mode: which refers to the case that the neuron generates a spike and means that the neuron's value is 1; 2) sleep mode: if the neuron's value is 0. Also, $w_{\cdot,\cdot}$ and $\mathbf{b}_{\cdot,\cdot}$ denote associated weights and biases for the network, respectively. Notice that $\mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \mathbf{g}_t$ can take the values 0, 1, or 2. Since the gradients around 2 are not as informative, we threshold this output to output 1 when it is 1 or 2. We approximate the gradients of this step function with γ that take two values 1 or ≤ 1 . Note that we can employ a Gaussian approximation at this step similar to our approach in the next section, and we observe that this relaxation does not affect the performance in practice, which is what we employed in the experiments.

3.2 Enabling Backpropagation in LSTM SNNs

Backpropagation is a major, if not the only, problem in SNNs. In this section, we proceed with an example. Regardless of the activations ($\sigma_1(u)$ or $\sigma_2(u)$), assume that we perturb the membrane potential of a neuron, $u_n(t)$, with an arbitrary random value δ_0 . Given $u_n(t)$, the neuron can be either in the wake mode or sleep mode. Based on the activation's threshold (see Figure 2b), this perturbation could switch the neuron's mode. For instance, in the wake mode if $\delta_0 < 0$ and also $u_n(t) + \delta_0 \leq \theta$ (θ is the threshold that can be either θ_1 or θ_2 based on the activation), the neuron will be forced to the sleep mode. With this, we can say that the change in neuron's mode is a function of the membrane potential and the threshold given by $|u| - |\theta|$. Therefore, if the mode switches the derivative of output w.r.t. $u_n(t)$ is proportional to $\sigma'(u + \delta_0) \propto \frac{\Delta\sigma(u)}{\Delta u} = \frac{1}{\delta_0}$, otherwise, $\sigma'(u + \delta_0) = 0$. Nevertheless, There is still a problem with small values of δ_0 that the mode switches (which equivalently means that $u_n(t)$ is close to the threshold). Indeed, this gradient will blow up the backpropagation of error.

To tackle with this issue, we suggest an alternative approximation. Consider the probability density function (pdf) $f(\delta)$ which corresponds to the pdf of changing mode with δ as the random variable. Given a small random perturbation δ_0 , the probability of switching mode is $\int_{\delta}^{\delta+\delta_0} f(\delta) \approx f(\delta + \delta_0)\delta_0$ and the probability of staying at the same mode is $1 - f(\delta + \delta_0)\delta_0$. As such, we can

Figure 2: (a)-Spiking neuron configuration; (b)-spike activation $\sigma_1(u)$; (c)-spike activations' derivatives $\sigma'_1(u)$ and $\sigma'_2(u)$; (d)-derivatives of spike $\sigma'(u)$ ($\alpha = 1$), sigmoid $S'(u)$, and $\tanh'(u)$ activations.

capture the expected value of $\sigma'(u)$ as follows:

$$\begin{aligned}
 \sigma'(u) &= \lim_{\delta_0 \rightarrow 0} E[\sigma'(u + \delta_0)] \\
 &= \lim_{\delta_0 \rightarrow 0} [f(\delta + \delta_0)\delta_0 \times \frac{1}{\delta_0} + (1 - f(\delta + \delta_0)\delta_0) \times 0] \\
 &= f(\delta) = f(|u| - |\theta|).
 \end{aligned} \tag{2}$$

It can be seen the activation's derivative could be relaxed with an appropriate symmetric (about the threshold θ) distribution, whose random variable δ is proportional to the difference neuron's membrane potential and the threshold, $|u| - |\theta|$.

We empirically observed that a good candidate for this distribution is the Gaussian distribution with suitable variance (see Figure 2c). Moreover, the smoothness of Gaussian distribution makes it a better candidate against other well-known symmetric distributions, i.e., Laplace distribution. Interestingly, another attribute that makes it unique is its curve which has, in spirit, analogous impact on backpropagation as the activations in traditional LSTM. In other words, Gaussian distribution has the same shape as the derivatives of the sigmoid and tanh activations. In addition, we can easily tune the variances corresponding to $\sigma'_1(u)$ and $\sigma'_2(u)$ to have the same shape as their counterpart activations in traditional LSTM (see Figure 2d).

3.3 Loss Function Derivative and Associated Parameter Updates

Next, we develop the update expressions for the parameters of LSTM spiking units. In order to do so, consider that the output layer is softmax, $\mathbf{y}_t = \text{softmax}(w\mathbf{y}_t + \mathbf{b}_y)$, and the loss function defined to be cross entropy loss. Therefore, the derivative of the loss function w.r.t. $\mathbf{y}_t$ output of LSTM SNNs at t can be characterized as follows:

$$\frac{\partial L}{\partial \mathbf{y}_t} = \mathbf{y}_t - \mathbf{y}_{\text{true}}, \quad (3)$$

where $\mathbf{y}_{\text{true}}$ is the true signal or label. Identically, networks with linear output layers and least square loss functions we have the same gradient. Given this and expressions in (1), the derivatives of the loss function w.r.t. outputs of each gate and layer can be derived as follows: All other derivatives with details are provided in Appendix A.

$$\begin{aligned} \frac{\partial L}{\partial \mathbf{h}_t} &= w_y \frac{\partial L}{\partial \mathbf{y}_t}, & \frac{\partial L}{\partial \mathbf{o}_t} &= \mathbf{c}_t \odot \frac{\partial L}{\partial \mathbf{h}_t}, & \frac{\partial L}{\partial \mathbf{c}_t} &= \gamma \mathbf{o}_t \odot \frac{\partial L}{\partial \mathbf{h}_t}, \\ \frac{\partial L}{\partial \mathbf{c}_{t-1}} &= \frac{\partial L}{\partial \mathbf{c}_{t-1}} + \mathbf{f}_t \odot \frac{\partial L}{\partial \mathbf{c}_t}, & \frac{\partial L}{\partial \mathbf{i}_t} &= \mathbf{g}_t \odot \frac{\partial L}{\partial \mathbf{c}_t}, \\ \frac{\partial L}{\partial \mathbf{g}_t} &= \mathbf{i}_t \odot \frac{\partial L}{\partial \mathbf{c}_t}, & \frac{\partial L}{\partial \mathbf{f}_t} &= \mathbf{c}_{t-1} \odot \frac{\partial L}{\partial \mathbf{c}_t}, \\ \frac{\partial L}{\partial \mathbf{h}_{t-1}} &= \sigma'_1(\Delta_1[w_{o,h}\mathbf{h}_{t-1} + w_{o,x}\mathbf{x}_t + \mathbf{b}_{o,h} + \mathbf{b}_{o,x}])w_{o,h} \frac{\partial L}{\partial \mathbf{o}_t} \\ &\quad + \sigma'_2(\Delta_2[w_{g,h}\mathbf{h}_{t-1} + w_{g,x}\mathbf{x}_t + \mathbf{b}_{g,h} + \mathbf{b}_{g,x}])w_{g,h} \frac{\partial L}{\partial \mathbf{g}_t} \\ &\quad + \sigma'_1(\Delta_1[w_{i,h}\mathbf{h}_{t-1} + w_{i,x}\mathbf{x}_t + \mathbf{b}_{i,h} + \mathbf{b}_{i,x}])w_{i,h} \frac{\partial L}{\partial \mathbf{i}_t} \\ &\quad + \sigma'_1(\Delta_1[w_{f,h}\mathbf{h}_{t-1} + w_{f,x}\mathbf{x}_t + \mathbf{b}_{f,h} + \mathbf{b}_{f,x}])w_{f,h} \frac{\partial L}{\partial \mathbf{f}_t} \end{aligned} \quad (4)$$

4 EXPERIMENTS

4.1 Settings and Datasets

We test our proposed method for different datasets. For all experiments, we initialize all weights based on standard normal distribution, and all biases are initialized to be zero at the beginning. Additionally, the networks are trained using Adam optimizer [18], with the learning rates of 0.001, $\beta_1 = 0.9$ and $\beta_2 = 0.999$ as the original paper. The thresholds for the spike activations have been set on $\theta_1, \theta_2 = 0.1$, which is optimized empirically. α_1 and α_2 are set to be 4 and 0.3, respectively. More details about this selection is provided in Appendix B.

4.2 Toy Dataset

We first illustrate the performance of the proposed method on a periodic sinusoidal signal. Our objective is to show that the proposed architecture can learn the temporal dependence using spikes as the input. Hence, we set our original input and target output to be $f(x) = 0.5 \sin(3x) + 0.5 \sin(6x) + 1$. In this case, the task is generating a prediction from a sequence of input spikes. To obtain this input spike train, after sampling the signal, we convert samples into ON- and OFF-event based values using Poisson process, where the value of each input shows the probability that it emits a spike as shown in Figure 3.

Next, we used the proposed deep LSTM spiking unit composed of one hidden layer of 100 spiking neurons and input size of 20. The output is a passed through a linear layer of size one. The loss function is $\frac{1}{2} \sum_{t=1}^T \|y_t - \hat{y}_t\|^2$, where y_t and $\hat{y}_t$ denote the actual and predicted outputs, respectively, and $T = 100$. Accordingly, we backpropagate the error using the proposed method. Also, we

Figure 3: Input signal and its spike representation after Poisson sampling. The value of input is assumed to be the probability that the associated neuron emits a spike.

empirically optimize α_1 and α_2 and set them to 4 and 0.3, respectively (more insight about the impacts of these parameters over the convergence rate and accuracy is provided for sequential MNIST dataset). The generated sequences and their convergence into true signal for different number of iterations are represented in Figure 4. As it shows, the network has learned the dependencies of samples in few iterations.

4.3 LSTM Spiking Network for Classification

Sequential MNIST [20] is a standard and popular dataset among machine learning researchers. The dataset consists of handwritten digits corresponding to 60k images for training and 10k test images. Each image is a 28×28 gray-scale pixels coming from 10 different classes. The main difference of sequential MNIST is that the network cannot get the whole image at once (see Appendix Figure 5). To convert each image to ON- and OFF-event based training samples we again use Poisson sampling, where the density of each pixel shows the probability that it emits a spike.

To make MNIST as a sequential dataset, we train the proposed LSTM spiking network over 28 time steps and input size of 28 for each time step (see Appendix Figure 5), and execute the optimization and let it run for 2000 epochs. Test accuracy and associated error bars are presented in Table 1. In addition, we listed the results of other state-of-the-art recurrent SNNs approaches for sequential MNIST, and feedforward SNNs for MNIST in the same table.

As it can be seen in Table 1, we achieve 98.3% test accuracy for sequential MNIST which is better than other LSTM-based SNNs and also this result is comparable to what was obtained by the feedforward SNN proposed in [17]. It should be noted that in [17] neurons are followed by time-based kernels and the network gets the whole image at once. Hereupon, We first note that kernel-based SNNs are not instantaneous. Usually, these networks are modeled continuously over time $t \in [0, T]$, and then are sampled with a proper sampling time T_s . For every time instance, each neuron goes through a convolution operation and finally the outputs are transferred to the next layer via matrix multiplication. This procedure is repeated for every time instance $t_s, s = 1, 2, \dots, \lceil \frac{1}{T_s} \rceil$. Even

Figure 4: Generated sequence for different number of iterations.

Table 1: Classification accuracy on sequential MNIST dataset. The conventional LSTM with the same architecture and parameters as this work gets 99.10% test accuracy.

METHOD	ARCHITECTURE	ACCURACY	BEST
CONVERTED FF SNN ^(A) [8]	784 – 500 – 500 – 10	94.09%	94.09%
FF SNN[17] ^(B)	784 – 800 – 10	98.93%	98.93%
LSTM SNN[5]	28 – 100 – 10 (28 LSTM UNITS)	97.29%	97.29%
LSTM SNN[1]	80 – 220 – 10 (128 LSTM UNITS)	96.4%	96.4%
THIS WORK	28 – 1000 – 10 (28 LSTM UNITS)	98.23 ± 0.07%	98.3%

(A) refers to indirect feedforward SNNs training, MNIST dataset,

(B) MNIST dataset,

though our proposed algorithm operates in discrete-time steps, one should note that the number of time steps in our model is much less compared to the kernel-based methods. Indeed, for kernel-based approaches one should prefer small sampling time to guarantee appropriate sampling, which, on the other hand, increases the number of time steps and consequently incurs more computation cost. For MNIST dataset, for example, the number of time steps required by our algorithm is 28 (see in Table 1), while the kernel-based method in [31] requires 350. Furthermore, in power-limited regimes computational complexity of kernel-based approaches make them less favorable candidates. However, in our proposed method, we eliminate the need for these kernels by drawing connections between LSTM and SNNs in order to model the dynamics of neurons. More information about the selection of α_1 and α_2 are provided in Appendix B.

Sequential EMNIST is another standard and relatively new benchmark for classification algorithms, which is an extended version of MNIST, but more challenging in the sense that it includes both letters and digits. It has almost 113K training, and about 19K test samples from 47 distinct classes. Using the same framework as sequential MNIST section, we convert the images into ON- and OFF- event-based sequential array for each image. Similarly, we train the network for 2000 iterations. The resulting test accuracy and the associated error bars are presented in Table. 2. The results of some other methods are also listed in the same table. Although this dataset has not been tested by other recurrent SNN approaches, we get comparable results with feedforward SNNs.

We believe there are several reasons for why FF SNN performs better in image classification tasks. Among them are getting the image at once, equipping each neuron with a time-based kernel and

sampling input multiple times (see [17] and [31]). However, RNNs in general and LSTM in particular have shown tremendous success in sequential learning tasks, which can be attributed to them equipping each neuron with an internal memory to manage the information flow from the sequential inputs. This feature leads RNN and its derivatives to be the preferred method in many sequential modeling tasks, especially in language modeling. FF networks, however, are not designed to learn the dependencies of a sequential input. While the proposed work in [17] performs better in image classification, it is not obvious how we can modify its architecture for sequential learning tasks, see the following experiments.

4.4 Language Modeling

The goal of this section is to demonstrate that the proposed LSTM SNN is also capable of learning high-quality language modeling tasks. By showing this, we can testify the network’s capability to learn long-term dependencies. In particular, we first train our network for prediction and then extend it to be a generative language model, for both character- and word-level, using the proposed LSTM SNN. Indeed, the proposed recurrent SNN will learn the dependencies in the strings of inputs and conditional probabilities of each character (word) given a sequence of characters (words). For both models, we use LSTM spiking unit with one hidden layer of size 200. Also, the same initialization and parameters as mentioned before.

Character-level - each dataset that we used for this part is a string of characters, including alphabets, digits, and punctuations. The network is a series of LSTM SNN units, and the input of each, $\mathbf{x}_t$, is a character which is one-hot encoded version of it, represented by the vector $(s_1, s_2, \dots, s_n)$, where n denote the total number of characters. Therefore, the input vector for each unit is a one-hot

Table 2: Classification accuracy on EMNIST dataset. The conventional LSTM with the same architecture and parameters as this work gets 87.1% test accuracy.

METHOD	ARCHITECTURE	ACCURACY	BEST
CONVERTED FF SNN [25]	784 – 200 – 200 – 47	81.77%	81.77%
FF SNN [25]	784 – 200 – 200 – 47	78.17%	78.17%
FF SNN [17]	784 – 800 – 47	85.41%	85.57%
THIS WORK (SEQUENTIAL EMNIST)	28 – 1000 – 47 (28 UNITS)	83.75 ± 0.15%	83.90%

Table 3: Test perplexity, character- and word-level

DATASET	CHARACTERS	LSTM SNN	LSTM	WORDS	LSTM SNN	LSTM
ALICE’S ADVENTURE	41	19.0267	14.7539	1.4K	85.3921	65.3658
WIKITEXT-2	74	19.3849	14.9319	2K	90.1725	86.1601

Table 4: Samples generated by LSTM SNN, *Alice’s Adventure in Wonderland* dataset

Generated Text (character-level)	Close Text
she is such a cring	she is such a nice
andone had no very clear notion all over with william	Alice had no very clear notion how long ago anything had happened
she began again: 'of hatting '	she began again: 'Ou est ma chatte?'
she was very like as thump!	she was not quite sure
Generated Text (word-level)	Close Text
alice began to get rather sleepy and went on	Alice began to feel very uneasy
the rabbit was no longer to be lost however on the second time round she could if i only knew how to begin	there was not a moment to be lost however on the second time round she came upon a low curtain

vector, which is also in favor of spike-based representation. Giving the training sequence $(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T)$, the network utilizes it to return the predictive sequence, denoted by $(o_1, o_2, \dots, o_T)$, where $o_{t+1} = \arg \max p(\mathbf{x}_{t+1} | \mathbf{x}_{\leq t})$. It should be noted that the last layer of each spiking LSTM module is a softmax.

The datasets that we employ are *Alice’s Adventures in Wonderland* and *Wikitext-2*. We first shrink these datasets and also clear them from capital letters by replacing small one. After this preprocessing, *Alice’s Adventures in Wonderland* and *Wikitext-2* include 41 and 74 distinct characters, respectively. And they both have 52000 total number of characters. Test dataset for this dataset is a different with the same distinct characters but total size of 12000. We used an LSTM with input size of characters, one hidden layer of size 200 and output size of characters. To evaluate the model, the averaged perplexity $(p(x_1, x_2, \dots, x_T)^{-1/T})$ after 1000 iterations is reported in Table 3. Also, we reported the results of conventional LSTM for similar datasets. As it can be seen the proposed LSTM spiking unit can achieve comparable results, however, its privilege is to be far more energy and resource efficient. After learning long-term dependencies successfully, the trained model also can be employed to generate text as well. Hence, to have a better vision of the quality

and richness of generated sequences, some samples are presented in Table 4.

Word-level - similar to character-level, we start by cleaning capital letters and then follow by extracting distinct words. However, compared to character-level, one-hot encoding for each word would be exhaustive. To tackle with this problem, we start by encoding each word to a representative vector. Based on word to vector, we use a window size of 5 (5 words behind and 5 word ahead), and train a feedforward neural network of one hidden layer with 100 units followed by a softmax layer. Hence, each word is represented by a vector of size 100, where different words with similar context are close to each other. In this representation, each vector carries critical information from words, and we expect significant loss of information when we convert vectors into spike-based representations. Therefore, we have input vectors in their main formats without any conversion to ON-and OFF-event based. Similar datasets to the previous task have been used. However, here we have an LSTM spiking unit with input size of 100, one hidden layer with 200 neurons, and output size of 100. Similar to the previous part, the results are provided in Table 3 and Table 4. Hence, for word-level language modeling task the results are also comparable with conventional LSTM.

4.5 Speech Classification

The goal of using speech recognition task is to evaluate the ability of our architecture to learn speech sequences for the classification task. To do so, we leverage a speech dataset recorded at 8kHz, FSDD, consisting of recordings of digits spoken from four different speakers, total size of 2000 (500 of each per speaker). To effectively represent each sample for training, first we transform samples using 1D wavelet scattering transform. After applying this pre-processing, each sample becomes a 1D vector size of 338 coming from 10 different classes. The proposed network for this task is a series of 8 LSTM spiking units, input size of 48 for each and the output is taken from the last unit where it is followed by a softmax layer. To evaluate the model, the dataset is divided into 1800 training and 200 test samples. Based on this methodology we achieved 86.3% accuracy for training set, and 83% for the test set. Employing the same architecture, training and test accuracy for conventional LSTM are 89.4% and 86%, respectively. It can be inferred that LSTM SNNs can get comparable results to conventional LSTM, but in a more efficient energy and resource manner.

5 CONCLUSION

In this work, we introduce a framework for direct-training of recurrent SNNs. In particular, we developed a class of LSTM-based SNNs that leverage the inherent LSTM capability of learning temporal dependencies. Based on this network, we develop a back-propagation framework for such networks based. We evaluate the performance of such LSTM SNNs over toy examples and then for the classification task. The results show that the proposed network achieve better performance compared to the existing recurrent SNNs. The results are also comparable with feedforward SNNs, while the proposed model is computationally less intensive. Finally, we test our method with a language modeling task to evaluate the performance of our network to learn long-term dependencies.

ACKNOWLEDGMENTS

This work was supported by ONR under grant N000141912590, and by ONR under grants 1731754 and 1564167.

REFERENCES

- [1] Guillaume Bellec et al. 2018. Long short-term memory and Learning-to-learn in networks of spiking neurons. In *Advances in Neural Information Processing Systems* 31. 787–797. <http://papers.nips.cc/paper/7359-long-short-term-memory-and-learning-to-learn-in-networks-of-spiking-neurons.pdf>
- [2] Sander M Bohte, Joost N Kok, and Han La Poutre. 2002. Error-backpropagation in temporally encoded networks of spiking neurons. *Neurocomputing* 48, 1-4 (2002), 17–37.
- [3] Olaf Booi and Hieu tat Nguyen. 2005. A gradient descent rule for spiking neurons emitting multiple spikes. *Inform. Process. Lett.* 95, 6 (2005), 552–558.
- [4] Gregory Cohen, Saeed Afshar, Jonathan Tapson, and André van Schaik. 2017. EMNIST: an extension of MNIST to handwritten letters. *arXiv preprint arXiv:1702.05373* (2017).
- [5] Rui Costa, Ioannis Alexandros Assael, Brendan Shillingford, Nando de Freitas, and Tim Vogels. 2017. Cortical microcircuits as gated-recurrent neural networks. In *Advances in Neural Information Processing Systems* 30. 272–283.
- [6] Mike Davies, Narayan Srinivasa, Tsung-Han Lin, Gautham Chinya, Yongqiang Cao, Sri Harsha Choday, Georgios Dimou, Prasad Joshi, Nabil Imam, Shweta Jain, et al. 2018. Loihi: A neuromorphic manycore processor with on-chip learning. *IEEE Micro* 38, 1 (2018), 82–99.
- [7] Gustavo Deco, Viktor K Jirsa, Peter A Robinson, Michael Breakspear, and Karl Friston. 2008. The dynamic brain: from spiking neurons to neural masses and cortical fields. *PLoS computational biology* 4, 8 (2008), e1000092.
- [8] Peter U Diehl, Daniel Neil, Jonathan Binas, Matthew Cook, Shih-Chii Liu, and Michael Pfeiffer. 2015. Fast-classifying, high-accuracy spiking deep networks through weight and threshold balancing. In *2015 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 1–8.
- [9] Steve K Esser et al. 2015. Backpropagation for Energy-Efficient Neuromorphic Computing. In *Advances in Neural Information Processing Systems* 28. 1117–1125.
- [10] Steve B Furber, David R Lester, Luis A Plana, Jim D Garside, Eustace Painkras, Steve Temple, and Andrew D Brown. 2012. Overview of the spinnaker system architecture. *IEEE Trans. Comput.* 62, 12 (2012), 2454–2467.
- [11] Wulfram Gerstner and Werner M Kistler. 2002. *Spiking neuron models: Single neurons, populations, plasticity*. Cambridge university press.
- [12] Samanwoy Ghosh-Dastidar and Hojjat Adeli. 2009. A new supervised learning algorithm for multiple spiking neural networks with application in epilepsy and seizure detection. *Neural networks* 22, 10 (2009), 1419–1431.
- [13] K. Greff et al. 2017. LSTM: A Search Space Odyssey. *IEEE Transactions on Neural Networks and Learning Systems* 28, 10 (Oct 2017), 2222–2232.
- [14] Robert Gütiğ and Haim Sompolsky. 2006. The tempotron: a neuron that learns spike timing-based decisions. *Nature neuroscience* 9, 3 (2006), 420.
- [15] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.
- [16] Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. 2016. Binarized Neural Networks. In *Advances in Neural Information Processing Systems* 29. Curran Associates, Inc., 4107–4115. <http://papers.nips.cc/paper/6573-binarized-neural-networks.pdf>
- [17] Yingyezhe Jin, Wenrui Zhang, and Peng Li. 2018. Hybrid Macro/Micro Level Backpropagation for Training Deep Spiking Neural Networks. In *Advances in Neural Information Processing Systems* 31. 7005–7015.
- [18] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [19] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*. 1097–1105.
- [20] Alex M Lamb, Anirudh Goyal ALIAS PARTH GOYAL, Ying Zhang, Saizheng Zhang, Aaron C Courville, and Yoshua Bengio. 2016. Professor Forcing: A New Algorithm for Training Recurrent Networks. In *Advances in Neural Information Processing Systems* 29. 4601–4609.
- [21] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *nature* 521, 7553 (2015), 436.
- [22] Qian Liu, Yunhua Chen, and Steve Furber. 2017. Noisy Softplus: an activation function that enables SNNs to be trained as ANNs. *arXiv preprint arXiv:1706.03609* (2017).
- [23] Sam McKennoch, Dingding Liu, and Linda G Bushnell. 2006. Fast Modifications of the SpikeProp Algorithm. In *IJCNN*, Vol. 6. 3970–3977.
- [24] Raoul-Martin Memmesheimer et al. 2014. Learning precisely timed spikes. *Neuron* 82, 4 (2014), 925–938.
- [25] Emre O Neftci, Charles Augustine, Somnath Paul, and Georgios Detorakis. 2017. Event-driven random back-propagation: Enabling neuromorphic deep learning machines. *Frontiers in neuroscience* 11 (2017), 324.
- [26] Peter O'Connor, Daniel Neil, Shih-Chii Liu, Tobi Delbruck, and Michael Pfeiffer. 2013. Real-time classification and sensor fusion with a spiking deep belief network. *Frontiers in neuroscience* 7 (2013), 178.
- [27] Ning Qiao, Hesham Mostafa, Federico Corradi, Marc Osswald, Fabio Stefanini, Dora Sumislawska, and Giacomo Indiveri. 2015. A reconfigurable on-line learning spiking neuromorphic processor comprising 256 neurons and 128K synapses. *Frontiers in neuroscience* 9 (2015), 141.
- [28] Johannes Schemmel, Daniel Brüderle, Andreas Gribbl, Matthias Hock, Karlheinz Meier, and Sebastian Millner. 2010. A wafer-scale neuromorphic hardware system for large-scale neural modeling. In *Proceedings of 2010 IEEE International Symposium on Circuits and Systems*. 1947–1950.
- [29] B. Schrauwen and J. Van Campenhout. 2004. Extending SpikeProp. In *2004 IEEE International Joint Conference on Neural Networks (IEEE Cat. No.04CH37541)*, Vol. 1. 471–475. <https://doi.org/10.1109/IJCNN.2004.1379954>
- [30] Amar Shrestha, Khadeer Ahmed, Yanzhi Wang, David P Widemann, Adam T Moody, Brian C Van Essen, and Qinru Qiu. 2017. A spike-based long short-term memory on a neuromorphic processor. In *2017 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*. IEEE, 631–637.
- [31] Sumit Bam Shrestha and Garrick Orchard. 2018. SLAYER: Spike Layer Error Reassignment in Time. In *Advances in Neural Information Processing Systems* 31. 1419–1428.
- [32] Sen Song, Kenneth D Miller, and Larry F Abbott. 2000. Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nature neuroscience* 3, 9 (2000), 919.
- [33] Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In *Advances in neural information processing systems*. 649–657.

Figure 5: (a)-28 LSTM spiking units for sequential MNIST classification task. Each row of the image is the input for the unit. Also, images are converted to their spike-representations using Poisson sampling; (b)-LSTM spiking units for language modeling task, character-level. Also, characters are represented to the unit using one-hot encoder.

A BACKPROPAGATION

We develop the update expressions for the parameters of LSTM spiking units. In order to do so, consider that the output layer is softmax, $\mathbf{y}_t = \text{softmax}(\mathbf{w}_y \mathbf{h}_t + \mathbf{b}_y)$, and the loss function defined to be cross entropy loss. Therefore, the derivative of the loss function w.r.t. $\mathbf{y}_t$ output of LSTM SNNs at t can be characterized as follows:

$$\frac{\partial L}{\partial \mathbf{y}_t} = \mathbf{y}_t - \mathbf{y}_{\text{true}}, \quad (5)$$

Identically, networks with linear output layers and least square loss functions we have the same gradient. Given this and also the derivatives of the loss function w.r.t. outputs of each gates in (4), we can now update the weights based on the derivative of the loss function for each of them:

$$\begin{aligned} d\mathbf{w}_y &= \sum_t \mathbf{h}_t^T \odot \frac{\partial L}{\partial \mathbf{y}_t}, \\ d\mathbf{w}_{o,x} &= \sum_t \sigma'_1(\Delta_1[\mathbf{w}_{o,h} \mathbf{h}_{t-1} + \mathbf{w}_{o,x} \mathbf{x}_t + \mathbf{b}_{o,h} + \mathbf{b}_{o,x}]) \mathbf{x}_t \frac{\partial L}{\partial \mathbf{o}_t}, \\ d\mathbf{w}_{o,h} &= \sum_t \sigma'_1(\Delta_1[\mathbf{w}_{o,h} \mathbf{h}_{t-1} + \mathbf{w}_{o,x} \mathbf{x}_t + \mathbf{b}_{o,h} + \mathbf{b}_{o,x}]) \mathbf{h}_{t-1} \frac{\partial L}{\partial \mathbf{o}_t}, \\ d\mathbf{w}_{i,x} &= \sum_t \sigma'_1(\Delta_1[\mathbf{w}_{i,h} \mathbf{h}_{t-1} + \mathbf{w}_{i,x} \mathbf{x}_t + \mathbf{b}_{i,h} + \mathbf{b}_{i,x}]) \mathbf{x}_{t-1} \frac{\partial L}{\partial \mathbf{i}_t}, \end{aligned}$$

Figure 6: Convergence rates for different values of α_1 and α_2 . Indeed, for $\alpha_1 = 4$, $\sigma'_1(u)$ have a similar curve as the derivative of the sigmoid activation $S'(u)$. Also, we can declare $\sigma'_2(u) \approx \tanh'(u)$ with $\alpha_1 = 0.3$.

$$\begin{aligned} d\mathbf{w}_{i,h} &= \sum_t \sigma'_1(\Delta_1[\mathbf{w}_{i,h} \mathbf{h}_{t-1} + \mathbf{w}_{i,x} \mathbf{x}_t + \mathbf{b}_{i,h} + \mathbf{b}_{i,x}]) \mathbf{h}_{t-1} \frac{\partial L}{\partial \mathbf{i}_t}, \\ d\mathbf{w}_{g,x} &= \sum_t \sigma'_2(\Delta_2[\mathbf{w}_{g,h} \mathbf{h}_{t-1} + \mathbf{w}_{g,x} \mathbf{x}_t + \mathbf{b}_{g,h} + \mathbf{b}_{g,x}]) \mathbf{x}_{t-1} \frac{\partial L}{\partial \mathbf{g}_t}, \\ d\mathbf{w}_{g,h} &= \sum_t \sigma'_2(\Delta_2[\mathbf{w}_{g,h} \mathbf{h}_{t-1} + \mathbf{w}_{g,x} \mathbf{x}_t + \mathbf{b}_{g,h} + \mathbf{b}_{g,x}]) \mathbf{h}_{t-1} \frac{\partial L}{\partial \mathbf{g}_t}, \\ d\mathbf{w}_{f,x} &= \sum_t \sigma'_1(\Delta_1[\mathbf{w}_{f,h} \mathbf{h}_{t-1} + \mathbf{w}_{f,x} \mathbf{x}_t + \mathbf{b}_{f,h} + \mathbf{b}_{f,x}]) \mathbf{x}_{t-1} \frac{\partial L}{\partial \mathbf{f}_t}, \\ d\mathbf{w}_{f,h} &= \sum_t \sigma'_1(\Delta_1[\mathbf{w}_{f,h} \mathbf{h}_{t-1} + \mathbf{w}_{f,x} \mathbf{x}_t + \mathbf{b}_{f,h} + \mathbf{b}_{f,x}]) \mathbf{h}_{t-1} \frac{\partial L}{\partial \mathbf{f}_t}, \end{aligned}$$

where $\Delta_1[\cdot] \triangleq |\cdot| - |\theta_1|$ and $\Delta_2[\cdot] \triangleq |\cdot| - |\theta_2|$, and γ is one or a positive number less than it (based on the value of $\mathbf{c}_t$, explained in 3.1 of the paper). Taking into account these partial derivatives at each time step t , we can now update the weights and biases based on the partial derivatives of the loss function with respect to them. And with same approach we can express the derivatives of the loss function for the biases.

B α_1 & α_2 IMPACTS

Figure 6 is depicted to reveal the serious effects of α_0 and α_1 on tuning the gradients. Indeed, these two parameters control the flow of error during the backpropagation for different parts of the LSTM spiking unit. An interesting point is that with $\alpha_1 = 4$ and $\alpha_2 = 0.3$, the LSTM SNNs becomes similar to conventional LSTM during the backpropagation. We have done these experiments on MNIST dataset. We observe the same outcomes for the other datasets as well.

Figure 7: (a)-SubLSTM unit [5]: all activations are replaced with the sigmoid, $S(u) = \frac{1}{1+e^{-u}}$, and also two multiplicative gates replaced with subtractive gates; (b)- LSTM spiking unit: all activations are replaced with the spike activations $\sigma_1(u)$ and $\sigma_2(u)$.