

TRAINING PROBABILISTIC SPIKING NEURAL NETWORKS WITH FIRST-TO-SPIKE DECODING

Alireza Bagheri[†], Osvaldo Simeone^{‡†}, and Bipin Rajendran[†]

[†]ECE Dept., New Jersey Institute of Technology, Newark, NJ 07102, USA.

Email: {ab745; bipin}@njit.edu

[‡]Dept. of Informatics, King's College London, London, WC2R 2LS, UK.

Email: osvaldo.simeone@kcl.ac.uk

ABSTRACT

Third-generation neural networks, or Spiking Neural Networks (SNNs), aim at harnessing the energy efficiency of spike-domain processing by building on computing elements that operate on, and exchange, spikes. In this paper, the problem of training a two-layer SNN is studied for the purpose of classification, under a Generalized Linear Model (GLM) probabilistic neural model that was previously considered within the computational neuroscience literature. Conventional classification rules for SNNs operate offline based on the number of output spikes at each output neuron. In contrast, a novel training method is proposed here for a first-to-spike decoding rule, whereby the SNN can perform an early classification decision once spike firing is detected at an output neuron. Numerical results bring insights into the optimal parameter selection for the GLM neuron and on the accuracy-complexity trade-off performance of conventional and first-to-spike decoding.

Index Terms— Spiking Neural Network (SNN), Generalized Linear Model (GLM), first-to-spike decoding, neuromorphic computing

1. INTRODUCTION

Most current machine learning methods rely on second-generation neural networks, which consist of simple static non-linear neurons. In contrast, neurons in the human brain are known to communicate by means of sparse spiking processes. As a result, they are mostly inactive, and energy is consumed sporadically. Third-generation neural networks, or Spiking Neural Networks (SNNs), aim at harnessing the energy efficiency of spike-domain processing by building on computing elements that operate on, and exchange, spikes [1]. SNNs can be natively implemented on neuromorphic chips that are currently being developed within academic projects and by major chip manufacturers. Proof-of-concept implementations have shown remarkable energy savings by multiple orders of magnitude with respect to second-generation neural networks (see, e.g., [2, 3]).

Notwithstanding the potential of SNNs, a significant stumbling block to their adoption is the dearth of flexible and

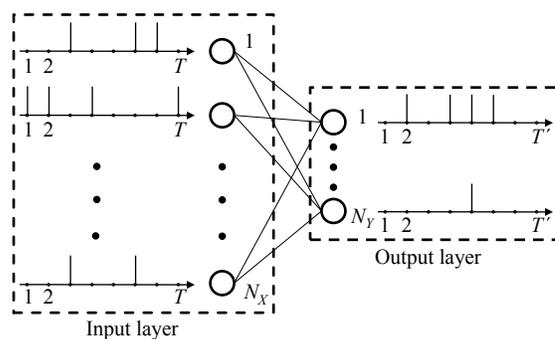


Fig. 1. Two-layer SNN for supervised learning.

effective learning algorithms. Most existing algorithms are based on variations of the unsupervised mechanism of Spike-Timing Dependent Plasticity (STDP), which updates synaptic weights based on local input and output spikes, and supervised variations that leverage global feedback [4, 5]. Another common approach is to convert trained second-generation networks to SNNs [6, 7]. Among the learning methods that attempt to directly maximize a spike-domain performance criterion, most techniques assume deterministic Spike Response Model (SRM) neurons, and propose various approximations to cope with the non-differentiability of the neurons' outputs (see [8, 9] and references therein).

While the use of probabilistic models for spiking neurons is standard in the context of computational neuroscience (see, e.g., [10]), probabilistic modeling has been sparsely considered in the machine learning literature on SNNs. This is despite the known increased flexibility and expressive power of probabilistic models [11, 12]. In the context of SNNs, as an example, probabilistic models have the capability of learning firing thresholds using standard gradient based methods, while in deterministic models these are instead treated as hyperparameters and set by using heuristic mechanisms such as homeostasis [13]. The state of the art on supervised learning with probabilistic models is set by [14] that considers Stochastic Gradient Descent (SGD) for Maximum Likelihood (ML) training, under the assumption that there exist given de-

sired output spike trains for all output neurons.

In this paper, we study the problem of training the two-layer SNN illustrated in Fig. 1 under a probabilistic neuron model, for the purpose of classification. Conventional decoding in SNNs operates offline by selecting the output neuron, and hence the corresponding class, with the largest number of output spikes [14]. In contrast, we study here a first-to-spike decoding rule, whereby the SNN can perform an early classification decision once a spike firing is detected at an output neuron. This generally reduces decision latency and complexity during the inference phase. The first-to-spike decision method has been investigated with temporal, rather than rate, coding and deterministic neurons in [15–18], but no learning algorithm exists under probabilistic neural models.

To fill this gap, we first propose the use of the flexible and computationally tractable Generalized Linear Model (GLM) that was introduced in [19] in the context of computational neuroscience (Section 3). Under this model, we then derive a novel SGD-based learning algorithm that maximizes the likelihood that the first spike is observed at the correct output neuron (Section 4). Finally, we present numerical results that bring insights into the optimal parameter selection for the GLM neuron and on the accuracy-complexity trade-off performance of conventional and first-to-spike decoding rules.

2. SPIKING NEURAL NETWORK WITH GLM NEURONS

In this section, we describe the architecture of the two-layer SNN under study and then we present the proposed GLM neuron model.

Architecture. We consider the problem of classification using a two-layer SNN. As shown in Fig. 1, the SNN is fully connected and has N_X presynaptic neurons in the input, or sensory layer, and N_Y neurons in the output layer. Each output neuron is associated with a class. In order to feed the SNN, an input example, e.g., a gray scale image, is converted to a set of N_X discrete-time spike trains, each with T samples, through rate encoding. The input spike trains are fed to the N_Y postsynaptic GLM neurons, which output discrete-time spike trains. A decoder then selects the image class on the basis of the spike trains emitted by the output neurons.

Rate encoding. With the conventional rate encoding method, each entry of the input signal, e.g., each pixel for images, is converted into a discrete-time spike train by generating an independent and identically distributed (i.i.d.) Bernoulli vectors. The probability of generating a “1”, i.e., a spike, is proportional to the value of the entry. In the experiments in Sec. 5, we use gray scale images with pixel intensities between 0 and 255 that yield a spike probability between 0 and 1/2.

GLM neuron model. The relationship between the input spike trains from the N_X presynaptic neurons and the output spike train of any postsynaptic neuron i follows a GLM, as illustrated in Fig. 2. To elaborate, we denote as $x_{j,t}$ and $y_{i,t}$ the binary signal emitted by the j -th presynaptic and the

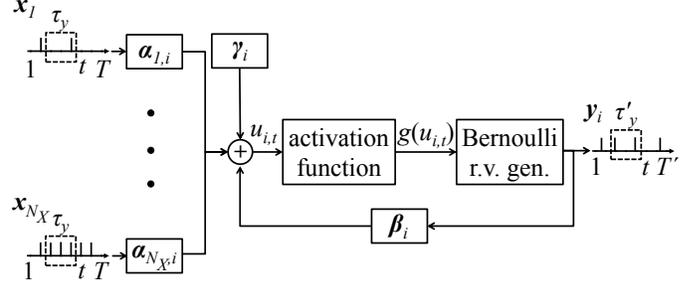


Fig. 2. GLM neuron model.

i -th postsynaptic neurons, respectively, at time t . Also, we let $\mathbf{x}_{j,a}^b = (x_{j,a}, \dots, x_{j,b})$ be the vector of samples from spiking process of the presynaptic neuron j in the time interval $[a, b]$. Similarly, the vector $\mathbf{y}_{i,a}^b = (y_{i,a}, \dots, y_{i,b})$ contains samples from the spiking process of the neuron i in the interval $[a, b]$. As seen in Fig. 2, the output $y_{i,t}$ of postsynaptic neuron i at time t is Bernoulli distributed, with firing probability that depends on the past spiking behaviors $\{\mathbf{x}_{j,t-\tau_y}^{t-1}\}$ of the presynaptic neurons $j = 1, \dots, N_X$ in a window of duration τ_y samples, as well as on the past spike timings $\mathbf{y}_{i,t-\tau'_y}^{t-1}$ of neuron i in a window of duration τ'_y samples. Mathematically, the membrane potential of postsynaptic neuron i at time t is given by

$$u_{i,t} = \sum_{j=1}^{N_X} \alpha_{j,i}^T \mathbf{x}_{j,t-\tau_y}^{t-1} + \beta_i^T \mathbf{y}_{i,t-\tau'_y}^{t-1} + \gamma_i, \quad (1)$$

where $\alpha_{j,i} \in \mathbb{R}^{\tau_y}$ is a vector that defines the *synaptic kernel* (SK) applied on the $\{j, i\}$ synapse between presynaptic neuron j and postsynaptic neuron i ; $\beta_i \in \mathbb{R}^{\tau'_y}$ is the *feedback kernel* (FK); and γ_i is a bias parameter. The vector of variable parameters θ_i includes the bias γ_i and the parameters that define the SK and FK filters, which are discussed below. Accordingly, the log-probability of the entire spike train $\mathbf{y}_i = [y_{i,1}, \dots, y_{i,T}]^T$ conditioned on the input spike trains $\mathbf{x} = \{\mathbf{x}_j\}_{j=1}^{N_X}$ can be written as

$$\log p_{\theta_i}(\mathbf{y}_i | \mathbf{x}) = \sum_{t=1}^T [y_{i,t} \log g(u_{i,t}) + \bar{y}_{i,t} \log \bar{g}(u_{i,t})], \quad (2)$$

where $g(\cdot)$ is an activation function, such as the sigmoid function $g(x) = \sigma(x) = 1/(1 + \exp(-x))$, and we defined $\bar{y}_{i,t} = 1 - y_{i,t}$ and $\bar{g}(u_{i,t}) = 1 - g(u_{i,t})$.

Unlike prior work on SNNs with GLM neurons, we adopt here the parameterized model introduced in [19] in the field of computational neuroscience. Accordingly, the SK and FK filters are parameterized as the sum of fixed basis functions with learnable weights. To elaborate, we write the SK $\alpha_{j,i}$ and the FK β_i as

$$\alpha_{j,i} = \mathbf{A} \mathbf{w}_{j,i}, \text{ and } \beta_i = \mathbf{B} \mathbf{v}_i, \quad (3)$$

respectively, where we have defined the matrices $\mathbf{A} = [\mathbf{a}_1, \dots, \mathbf{a}_{K_\alpha}]$ and $\mathbf{B} = [\mathbf{b}_1, \dots, \mathbf{b}_{K_\beta}]$ and the vectors $\mathbf{w}_{j,i} = [w_{j,i,1}, \dots, w_{j,i,K_\alpha}]^T$ and $\mathbf{v}_i = [v_{i,1}, \dots, v_{i,K_\beta}]^T$; K_α and K_β denote the respective number of basis functions; $\mathbf{a}_k =$

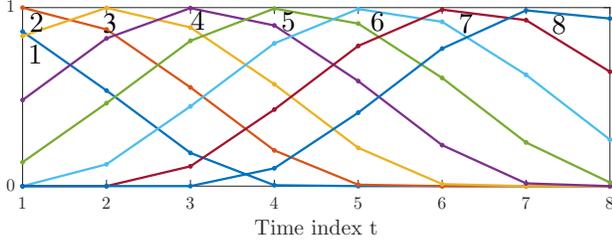


Fig. 3. Basis functions used in Sec. 5 ($a = 7500$ and $c = 1$ in [19, Sec. Methods]).

$[a_{k,1}, \dots, a_{k,\tau_y}]^T$ and $\mathbf{b}_k = [b_{k,1}, \dots, b_{k,\tau'_y}]^T$ are the basis vectors; and $\{w_{j,i,k}\}$ and $\{v_{i,k}\}$ are the learnable weights for the kernels $\alpha_{j,i}$ and β_i , respectively. This parameterization generalizes previously studied models for machine learning application. For instance, as a special case, if we set $K_\alpha = K_\beta = 1$, set \mathbf{a}_1 and \mathbf{b}_1 as in [14, eqs. (4) and (5)], and fix the weights $v_{i,1} = 1$, equation (2) yields a discrete-time approximation of the model considered in [14]. As another example, if we set $K_\alpha = \tau_y$, $K_\beta = \tau'_y$, $\mathbf{a}_k = \mathbf{1}_k$, $\mathbf{b}_k = \mathbf{1}_k$, where $\mathbf{1}_k$ is the all-zero vector except for a one in position k , (1) yields the unstructured GLM model considered in [20]. For the experiments discussed in Sec. 5, we adopt the time-localized raised cosine basis functions introduced in [19], which are illustrated in Fig. 3. Note that this model is flexible enough to include the learning of synaptic delays [21, 22].

3. TRAINING WITH CONVENTIONAL DECODING

In this section, we briefly review ML training based on conventional rate decoding for the two-layer SNN. During the inference phase, decoding is conventionally carried out in post-processing by selecting the output neuron with the largest number of spikes. In order to facilitate the success of this decoding rule, in the training phase, the postsynaptic neuron corresponding to the correct label $c \in \{1, \dots, N_Y\}$ is typically assigned a desired output spike train \mathbf{y}_c with a number of spikes, while a zero output is assigned to the other postsynaptic neurons \mathbf{y}_i with $i \neq c$.

Using the ML criterion, one hence maximizes the sum of the log-probabilities (2) of the desired output spikes $\mathbf{y}(c) = \{\mathbf{y}_1(c), \dots, \mathbf{y}_{N_Y}(c)\}$ for the correct label c given the N_X input spike trains $\mathbf{x} = \{\mathbf{x}_1, \dots, \mathbf{x}_{N_X}\}$, i.e., $L(\boldsymbol{\theta}) = \sum_{i=1}^{N_Y} \log p_{\theta_i}(\mathbf{y}_i(c) | \mathbf{x})$. The sum is further extended to all examples in the training set. The parameter vector $\boldsymbol{\theta} = \{\mathbf{W}, \mathbf{V}, \boldsymbol{\gamma}\}$ includes the parameters $\mathbf{W} = \{\mathbf{W}_i\}_{i=1}^{N_Y}$, $\mathbf{V} = \{\mathbf{v}_i\}_{i=1}^{N_Y}$ and $\boldsymbol{\gamma} = \{\gamma_i\}_{i=1}^{N_Y}$. The negative log-likelihood $-L(\boldsymbol{\theta})$ is convex with respect to $\boldsymbol{\theta}$ and can be minimized via SGD.

4. TRAINING WITH FIRST-TO-SPIKE DECODING

In this section, we introduce the proposed learning approach based on GLM neurons and first-to-spike decoding.

During the inference phase, with first-to-spike decoding, a decision is made once a first spike is observed at an output neuron. In order to train the SNN for this classification

rule, we propose to follow the ML criterion by maximizing the probability to have the first spike at the output neuron corresponding to the correct label c . The logarithm of this probability for a given example \mathbf{x} can be written as

$$L(\boldsymbol{\theta}) = \log \left(\sum_{t=1}^T p_t(\boldsymbol{\theta}) \right), \quad (4)$$

where

$$p_t(\boldsymbol{\theta}) = \prod_{i=1, i \neq c}^{N_Y} \prod_{t'=1}^t \bar{g}(u_{i,t'}) g(u_{c,t}) \prod_{t'=1}^{t-1} \bar{g}(u_{c,t'}), \quad (5)$$

is the probability of having the first spike at the correct neuron c at time t . In (5), the potential $u_{i,t}$ for all i is obtained from (1) by setting $y_{i,t} = 0$ for all i and t . The log-likelihood function $L(\boldsymbol{\theta})$ in (4) is not concave, and we tackle its maximization via SGD.

To this end, the gradients of the log-likelihood function for a given input \mathbf{x} can be computed after some algebra as (see [23] for details)

$$\begin{aligned} \nabla_{\mathbf{w}_{j,i}} L(\boldsymbol{\theta}) &= \begin{cases} - \sum_{t=1}^T \rho_{i,t} h_t g(u_{i,t}) \mathbf{A}^T \mathbf{x}_{j,t-\tau_y}^{t-1} & i \neq c \\ - \sum_{t=1}^T \rho_{c,t} (h_t g(u_{c,t}) - q_t) \mathbf{A}^T \mathbf{x}_{j,t-\tau_y}^{t-1} & i = c \end{cases}, \end{aligned} \quad (6)$$

for the weights and

$$\nabla_{\gamma_i} L(\boldsymbol{\theta}) = \begin{cases} - \sum_{t=1}^T \rho_{i,t} h_t g(u_{i,t}) & i \neq c \\ - \sum_{t=1}^T \rho_{c,t} (h_t g(u_{c,t}) - q_t) & i = c \end{cases}, \quad (7)$$

for the biases, where we have defined

$$\rho_{i,t} = \frac{g'(u_{i,t})}{g(u_{i,t}) \bar{g}(u_{i,t})}, \quad (8)$$

and

$$h_t = \sum_{t'=t}^T q_{t'} = 1 - \sum_{t'=1}^{t-1} q_{t'}, \quad (9)$$

with

$$q_t = \frac{p_t(\boldsymbol{\theta})}{\sum_{t'=1}^T p_{t'}(\boldsymbol{\theta})}. \quad (10)$$

Note that we have $\rho_{i,t} = 1$ when g is the sigmoid function.

Based on (6), the resulting SGD update can be considered as a neo-Hebbian rule [24], since it multiplies the contributions of the presynaptic neurons and of the postsynaptic activity, where the former depends on \mathbf{x} and the latter on the potential $u_{i,t}$. Furthermore, in (6)–(7), the probabilities $g(u_{i,t})$ and $g(u_{c,t})$ of firing at time t are weighted by the probability h_t in (9). By (10), this is the probability that the correct neuron is the first to spike and that it fires at some time $t' \geq t$, given that it is the first to spike at some time in the interval $[1, 2, \dots, T]$.

As a practical note, in order to avoid vanishing values in calculating the weights (10), we compute each probability term $p_t(\boldsymbol{\theta})$ in the log-domain, and normalize all the resulting terms with respect to the minimum proba-

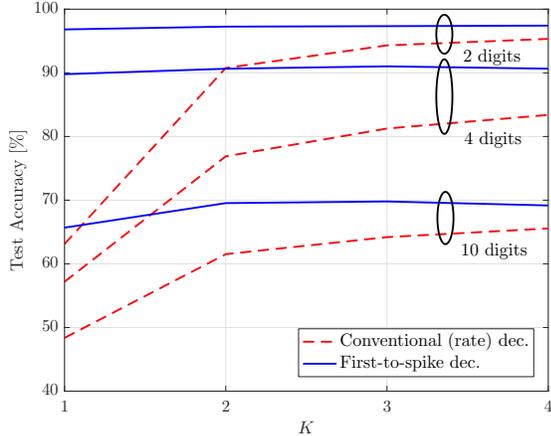


Fig. 4. Test accuracy versus the number K of basis functions for both conventional (rate) and first-to-spike decoding rules when $T = 4$.

bility as $q_t = \exp(a_t) / \sum_{t'=1}^T \exp(a_{t'})$, where $a_t = \ln(p_t) - \min_t(\ln(p_t))$.

5. NUMERICAL RESULTS

In this section, we numerically study the performance of the probabilistic SNN in Fig. 1 under conventional and first-to-spike decoding rules. We use the standard MNIST dataset [25] as the input data. As a result, we have $N_X = 784$, with one input neuron per pixel of the \mathbf{x} images. Following [3], we consider different number of classes, or digits, namely, the two digits $\{5, 7\}$, the four digits $\{5, 7, 1, 9\}$ and all 10 digits $\{0, \dots, 9\}$, and we use 1000 samples of each class for training and the same number for test set. We use a desired spike train with one spike after every three zeros for training the conventional decoding. SGD with minibatch size of one with 200 training epochs is used for both schemes. Ten-fold cross-validation is applied for selecting between 10^{-3} or 10^{-4} for the constant learning rates. The model parameters θ are randomly initialized with uniform distribution between -1 and 1.

We evaluate the performance of the schemes in terms of classification accuracy in the test set and of inference complexity. The inference complexity is measured by the total number of elementary operations, namely additions and multiplications, for input image that are required by the SNN during inference. The number of arithmetic operations needed to calculate the membrane potential (1) of neuron i at time instant t is of the order of $\mathcal{O}(N_X \tau_y + \tau'_y)$. As a result, in the conventional decoding method, the inference complexity per output neuron, or per class, is of the order $\mathcal{O}(T(N_X s_x + s_y))$, where s_x and s_y are the fraction of spikes in \mathbf{x} and \mathbf{y} , respectively. In contrast, with the first-to-spike decoding rule, the SNN can perform an early decision once a single spike is detected, and hence its complexity order is $\mathcal{O}(t(N_X s_x + s_y))$, where $1 \leq t \leq T$ is the (random) decision time.

We first consider the test classification accuracy as a function of the number K of basis functions in the GLM neural

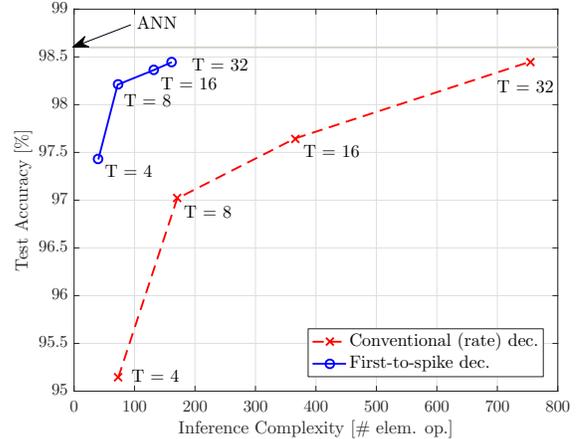


Fig. 5. Test accuracy versus per-class inference complexity for both conventional (rate) and first-to-spike decoding rules.

model. The basis functions are numbered as in Fig. 3, and we set $T = 4$. From Fig. 4, we observe that conventional decoding requires a large number K in order to obtain its best accuracy. This is due to the need to ensure that the correct output neuron fires consistently more than the other neurons in response to the input spikes. This, in turn, requires a larger temporal reception field, i.e., a larger K , to be sensitive to the randomly located input spikes. We note that for small values of T , such as $T = 4$, first-to-spike decoding obtains better accuracies than conventional decoding.

Fig. 5 depicts the test classification accuracy versus the inference complexity for both conventional and first-to-spike decoding rules for two digits when $K = T$. The classification accuracy of a conventional two-layer artificial neural network (ANN) with logistic neurons is added for comparison. From the figure, first-to-spike decoding is seen to offer a significantly lower inference complexity, thanks to its capability for early decisions, without compromising the accuracy. For instance, when the classification accuracy equals to 98.4%, the complexity of the conventional decoding method is five times larger than the first-to-spike method. Note also that conventional decoding generally requires large values of T to perform satisfactorily.

6. CONCLUSIONS

In this paper, we have proposed a novel learning method for probabilistic two-layer SNN that operates according to the first-to-spike learning rule. We have demonstrated that the proposed method improves the accuracy-inference complexity trade-off with respect to conventional decoding. Additional work is needed in order to generalize the results to multi-layer networks.

7. ACKNOWLEDGMENT

This work was supported by the U.S. NSF under grant ECCS #1710009. O. Simeone has also received funding from the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation program (grant agreement #725731).

8. REFERENCES

- [1] H. Paugam-Moisy and S. Bohte, "Computing with spiking neuron networks," in *Handbook of natural computing*, 2012, pp. 335–376.
- [2] J. Vincent, "Intel investigates chips designed like your brain to turn the AI tide," <https://www.theverge.com/2017/9/26/16365390/intel-investigates-chips-designed-like-your-brain-to-turn-the-ai-tide>, accessed: Sept. 26, 2017.
- [3] A. Diamond, T. Nowotny, and M. Schmukey, "Comparing neuromorphic solutions in action: implementing a bio-inspired solution to a benchmark classification task on three parallel-computing platforms," *Front. Neurosci.*, vol. 9, p. 491, 2016.
- [4] F. Ponulak and A. Kasiński, "Supervised learning in spiking neural networks with ReSuMe: sequence learning, classification, and spike shifting," *Neural Comput.*, vol. 22, no. 2, pp. 467–510, 2010.
- [5] R. V. Florian, "Reinforcement learning through modulation of spike-timing-dependent synaptic plasticity," *Neural Comput.*, vol. 19, no. 6, pp. 1468–1502, 2007.
- [6] P. O'Connor and M. Welling, "Deep spiking networks," *arXiv preprint arXiv:1602.08323*, 2016.
- [7] E. Hunsberger and C. Eliasmith, "Spiking deep networks with LIF neurons," *arXiv preprint arXiv:1510.08829*, 2015.
- [8] N. Anwani and B. Rajendran, "NormAD-normalized approximate descent based supervised learning rule for spiking neurons," in *Proc. IEEE Int. Joint Conf. on Neural Netw. (IJCNN)*, 2015.
- [9] J. H. Lee, T. Delbruck, and M. Pfeiffer, "Training deep spiking neural networks using backpropagation," *Front. Neurosci.*, vol. 10, p. 508, 2016.
- [10] J. W. Pillow, L. Paninski, V. J. Uzzell, E. P. Simoncelli, and E. J. Chichilnisky, "Prediction and decoding of retinal ganglion cell responses with a probabilistic spiking model," *J. Neurosci.*, vol. 25, no. 47, pp. 11 003–11 013, 2005.
- [11] D. Koller and N. Friedman, *Probabilistic graphical models: principles and techniques*. MIT press, 2009.
- [12] O. Simeone, "A brief introduction to machine learning for engineers," *arXiv preprint arXiv:1709.02840*, 2017.
- [13] R. Jolivet, A. Rauch, H. Lüscher, and W. Gerstner, "Predicting spike timing of neocortical pyramidal neurons by simple threshold models," *J. Comput. Neurosci.*, vol. 21, no. 1, pp. 35–49, 2006.
- [14] B. Gardner and A. Grüning, "Supervised learning in spiking neural networks for precise temporal encoding," *PLoS ONE*, vol. 11, no. 8, p. e0161335, 2016.
- [15] T. Masquelier and S. J. Thorpe, "Unsupervised learning of visual features through spike timing dependent plasticity," *PLoS Comput. Biol.*, vol. 3, no. 2, p. e31, 2007.
- [16] M. Mozafari, S. Kheradpisheh, T. Masquelier, A. Nowzari-Dalini, and M. Ganjtabesh, "First-spike based visual categorization using reward-modulated STDP," *arXiv preprint arXiv:1705.09132*, 2017.
- [17] J. Wang, A. Belatreche, L. P. Maguire, and T. M. McGinnity, "SpikeTemp: An enhanced rank-order-based learning approach for spiking neural networks with adaptive structure," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 28, no. 1, pp. 30–43, 2017.
- [18] Z. Lin, D. Ma, J. Meng, and L. Chen, "Relative ordering learning in spiking neural network for pattern recognition," *Neurocomputing*, vol. 275, pp. 94–106, 2018.
- [19] J. W. Pillow, J. Shlens, L. Paninski, A. Sher, A. M. Litke, E. J. Chichilnisky, and E. P. Simoncelli, "Spatio-temporal correlations and visual signalling in a complete neuronal population," *Nature*, vol. 454, no. 7207, pp. 995–999, 2008.
- [20] J. Shlens, "Notes on generalized linear models of neurons," *arXiv preprint arXiv:1404.1999*, 2014.
- [21] P. Baldi and A. F. Atiya, "How delays affect neural dynamics and learning," *IEEE Transactions on Neural Networks*, vol. 5, no. 4, pp. 612–621, 1994.
- [22] A. Taherkhani, A. Belatreche, Y. Li, and L. P. Maguire, "DL-ReSuMe: a delay learning-based remote supervised method for spiking neurons," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 26, no. 12, pp. 3137–3149, 2015.
- [23] A. Bagheri, O. Simeone, and B. Rajendran, "Training probabilistic spiking neural networks with first-to-spike decoding," *arXiv preprint arXiv:1710.10704*, 2017.
- [24] N. Frémaux and W. Gerstner, "Neuromodulated spike-timing-dependent plasticity, and theory of three-factor learning rules," *Front. Neural Circuits*, vol. 9, 2015.
- [25] Y. LeCun, "The MNIST database of handwritten digits," <http://yann.lecun.com/exdb/mnist/>, 1998.