Training Energy-Efficient Deep Spiking Neural Networks with Single-Spike Hybrid Input Encoding

Gourav Datta, Souvik Kundu, Peter A. Beerel
Ming Hsieh Department of Electrical and Computer Engineering
University of Southern California
Los Angeles, California 90089, USA
{gdatta, souvikku, pabeerel}@usc.edu

Abstract—**Spiking Neural Networks (SNNs) have emerged as an attractive alternative to traditional deep learning frameworks, since they provide higher computational efficiency in event driven neuromorphic hardware. However, the state-of-the-art (SOTA) SNNs suffer from high inference latency, resulting from inefficient input encoding and training techniques. The most widely used input coding schemes, such as Poisson based rate-coding, do not leverage the temporal learning capabilities of SNNs. This paper presents a training framework for low-latency energy-efficient SNNs that uses a hybrid encoding scheme at the input layer in which the analog pixel values of an image are directly applied during the first timestep and a novel variant of spike temporal coding is used during subsequent timesteps. In particular, neurons in every hidden layer are restricted to fire at most once per image which increases activation sparsity. To train these hybrid-encoded SNNs, we propose a variant of the gradient descent based spike timing dependent backpropagation (STDB) mechanism using a novel cross entropy loss function based on both the output neurons' spike time and membrane potential. The resulting SNNs have reduced latency and high activation sparsity, yielding significant improvements in computational efficiency. In particular, we evaluate our proposed training scheme on image classification tasks from CIFAR-10 and CIFAR-100 datasets on several VGG architectures. We achieve top-1 accuracy of 66.46% with 5 timesteps on the CIFAR-100 dataset with $\sim 125\times$ less compute energy than an equivalent standard ANN. Additionally, our proposed SNN performs $5\text{-}300\times$ faster inference compared to other state-of-the-art rate or temporally coded SNN models.**

Index Terms—**SNN, STDB, Input encoding, Energy-efficient SNNs**

I. INTRODUCTION

Artificial Neural Networks (ANNs) have contributed to a number of impressive success stories in Artificial General Intelligence (AGI) [1]–[5]. However, their superior performance has come at the cost of high computational and memory requirements [6], [7]. While convolutional neural networks (CNNs) on general purpose high-performance compute platforms such as GPUs are now ubiquitous [8], there has been increasing interest in domain-specific hardware accelerators [9] and alternate types of neural networks. In particular, Spiking Neural Network (SNN) accelerators have emerged as a potential low power alternative for AGI [10]–[13]. SNNs attempt to emulate the remarkable energy-efficiency of the brain with event-driven neuromorphic hardware. Neurons in an SNN exchange information via discrete binary events, resulting in a significant paradigm shift from traditional CNNs.

Because SNNs receive and transmit information through spikes, analog values must be encoded into a sequence of spikes. There has been a plethora of encoding methods proposed, including rate coding [14], [15], temporal coding [16]–[19], rank-order coding [20], phase coding [21], [22] and other exotic coding schemes [23]. Among these, rate-coding has shown competitive performance on complex tasks [14], [15] while others are either generally limited to simple tasks

such as learning the XOR function and classifying digits from the MNIST dataset or require a large number of spikes for inference. In rate coding, the analog value is converted to a spike train using a Poisson generator function with a rate proportional to the input pixel value. The number of timesteps in each train is inversely proportional to the quantization error in the representation, as illustrated in Fig. 1(b). Low error requirements force a large number of timesteps at the expense of high inference latency and low activation sparsity [15]. Temporal coding, on the other hand, has higher sparsity and can more explicitly represent correlations in inputs. However, temporal coding is challenging to scale [20] to vision tasks and often requires kernel-based spike response models [17] which are computationally expensive compared to the traditional leaky-integrate-and-fire (LIF) or integrate-and-fire (IF) models. Recently, the authors in [24] proposed direct input encoding, where they feed the analog pixel values directly into the first convolutional layer, which treats them as input currents to LIF neurons. Another recently proposed temporal encoding scheme uses the discrete cosine transform (DCT) to distribute the spatial pixel information over time for learning low-latency SNNs [25]. However, up to now, there has been no attempt to combine both spatial (captured by rate or direct encoding) and temporal information processed by the SNNs.

In addition to accommodating the various of forms of encoding inputs, supervised learning algorithms for SNNs have overcome many roadblocks associated with the discontinuous derivative of the spike activation function [26]–[28]. However, effective SNN training remains a challenge, as seen by the fact that SNNs still lag behind ANNs in terms of latency and accuracy in traditional classification tasks [15], [29]. A single feed-forward pass in ANN corresponds to multiple forward passes in SNN which is associated with a fixed number of timesteps. In spike-based backpropagation, the backward pass requires the gradients to be integrated over every timestep which increases computation and memory complexity [26], [30]. It requires multiple iterations, is memory intensive (for backward pass computations), and energy-inefficient, and thus has been mainly limited to small datasets (e.g. CIFAR-10) on simple shallow convolutional architectures [30]. Researchers have also observed high spiking activity and energy consumption in these trained SNN models [31], which further hinders their deployment in edge applications. Thus, the current challenges in SNN models are high inference latency and spiking activity, long training time, and high training costs in terms of memory and computation.

To address these challenges, this paper makes the following contributions:

- *Hybrid Spatio-Temporal Encoding:* We employ a hybrid input encoding technique where the real-valued image pixels are fed to the SNN during the first timestep. During

the subsequent timesteps, the SNN follows a single-spike temporal coding scheme, where the arrival time of the input spike is inversely proportional to the pixel intensity. While the direct encoding in the first timestep helps the SNN achieve low inference latency, the temporal encoding increases activation sparsity.

- *Single Spike LIF Model:* To further harness the benefits of temporal coding, we propose a modified LIF model, where neurons in every hidden layer fire at most once over all the timesteps. This leads to higher activation sparsity and compute efficiency.
- *Novel Loss Function:* We also propose a variant of the gradient descent based spike timing dependent backpropagation mechanism to train SNNs with our proposed encoding technique. In particular, we employ a hybrid cross entropy loss function to capture both the accumulated membrane potential and the spike time of the output neurons.

The remainder of our paper is structured as follows. In Section II we present the necessary background. Section III describes our proposed input encoding technique. We present our detailed experimental evaluation of the classification accuracy and latency in Section V. We show the energy improvement of our proposed framework in Section VI and finally present conclusions in Section VII.

II. BACKGROUND

A. SNN Fundamentals

An SNN consists of a network of neurons that communicate through a sequence of spikes modulated by synaptic weights. The spiking dynamics of a neuron are typically represented using either Integrate-and-Fire (IF) [32] or Leaky-Integrate-and-Fire (LIF) model [33]. Fig. 1(a) illustrates a basic SNN architecture with IF neurons processing rate-coded inputs. Both IF and LIF neurons integrate the input current into their respective states referred to as membrane potentials. The key difference between the models is that the membrane potential of a IF neuron does not change during the time period between successive input spikes while the LIF neuronal membrane potential leaks with a finite time constant. In this work, we use the LIF model to convert ANNs trained with ReLU activations to SNNs, because the leaky behaviour provides improved robustness to noisy spike-inputs and better generalization compared to those with no leak [34]. Moreover, the leak term provides a tunable control knob, which can be leveraged to improve inference accuracy, latency, and spiking activity in SNNs.

To characterize the LIF model, we use the following differential equation

$$C\frac{dU_i^t}{dt} + GU_i^t = I_i^t = \sum_j W_{ij} \cdot S_j^t \qquad (1)$$

where C and G are the membrane capacitance and conductance respectively. U_i^t and I_i^t are the membrane potential and input synaptic current of the i^{th} neuron at time t. Note that U_i^t integrates the incoming (pre-neuron) spikes S_j^t modulated by weights W_{ij} and leaks with a time constant equal to $\frac{C}{G}$. The post-neuron generates an output spike when U_i exceeds the firing threshold V. However, because of its' continuous representation, Eq. 1 is not suitable for implementations in popular Machine Learning (ML) frameworks (eg. Pytorch). Hence, we convert Eq. 1 into an iterative discrete-time version,
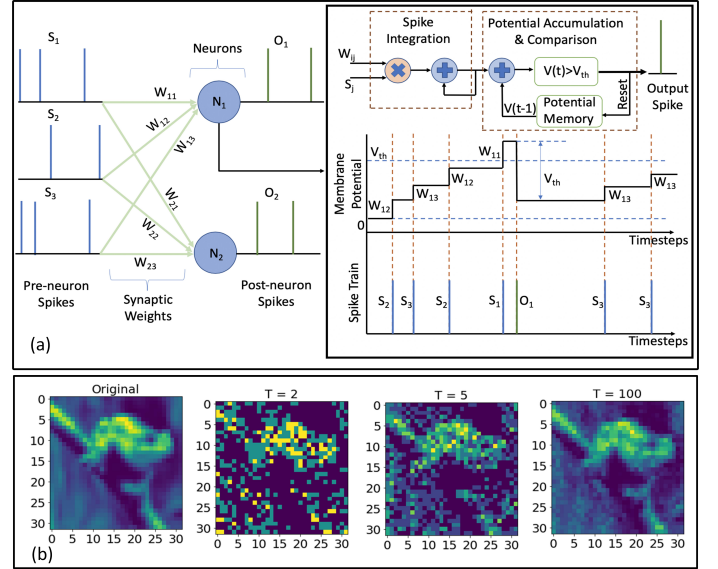

Fig. 1. (a) Feedforward fully-connected SNN architecture with Integrate and Fire (IF) spiking dynamics, (b) The spike input generated over several timesteps through Poisson generator. It is clear that having more timesteps yields a better approximation of the input image.

as shown in Eq. 2 [30], in which spikes are characterized as binary values (1 represents the presence of a spike). Note that λ represents the leak term which reduces U_i by a factor of $(1-\lambda)$ in every timestep.

$$U_i^t = \lambda U_i^{t-1} + \sum_j W_{ij} S_j(t) - VO_i^{t-1} \qquad (2)$$

The binary output spike at timestep t is given as

$$O_i^t = \begin{cases} 1, & \text{if } U_i^t > V \\ 0, & \text{otherwise} \end{cases} \qquad (3)$$

Note that the last term in Eq. 2 represents soft reset that reduces the membrane potential U_i by the threshold V at timestep t in response to an output spike generated at timestep $(t-1)$. In contrast, hard reset means resetting U_i to 0 after an output spike is generated. Soft reset minimises the information loss by allowing the spiking neuron to carry forward the surplus potential above the firing threshold to the subsequent timestep [30], [35] and is adopted in this work.

B. SNN Training Techniques

Recent research on training supervised deep SNNs can be broadly divided into three categories: i) Indirect learning; ii) Direct Learning; iii) Hybrid Learning.

1) Indirect Learning: Recent works have demonstrated that SNNs can be efficiently converted from ANNs by approximating the activation value of ReLU neurons with the firing rate of spiking neurons [12], [15], [36]–[38]. This technique uses the standard backpropagation algorithm for training in the ANN domain, and helps SNNs achieve SOTA results on various challenging inference tasks, particularly in image recognition [15], [36]. Moreover, ANN-SNN conversion simplifies the training procedures compared to approximate gradient techniques, since it involves only a single forward pass to process a single input. However, a disadvantage of ANN-SNN conversion is that it yields SNNs with an order of magnitude higher latency than other training techniques [15]. In this work,

we use ANN-SNN conversion as an initial step in our proposed framework because it yields high classification accuracy on deep networks. We then leverage direct encoding in the first timestep to reduce the number of synaptic operations and thus improve the SNN's energy efficiency.

2) Direct Learning: The discontinuous and non-differentiable nature of a spiking neuron makes it difficult to implement gradient descent based backpropagation. Consequently, several approximate training methodologies have been proposed that leverage the temporal dynamics of SNNs [26], [39]–[43]. The basic idea of these works is to approximate the spiking neuron functionality with a continuous differentiable model or use surrogate gradients to approximate real gradients. However, STDB requires the gradients to be integrated over all timesteps, increasing computation and memory requirements significantly, particularly for deep networks.

3) Hybrid Learning: Authors in [30] proposed a hybrid training methodology that consists of ANN-SNN conversion, followed by approximate gradient descent on the initialized network to obtain the final trained SNN model. The authors claimed that combining the two training techniques helps SNNs converge within a few epochs and require fewer timesteps. Another recent paper [24] proposes a training scheme for deep SNNs in which the membrane leak and the firing threshold along with other network parameters (weights) are updated at the end of every batch via gradient descent after ANN-SNN conversion. Moreover, instead of converting the image pixel values into spike trains using Poisson rate coding described above, the authors directly feed the analog pixel values in the first convolutional layer, which emits spikes using the LIF neuron model. This enables requiring fewer timesteps compared to Poisson rate coding. In this work, we employ a variant of the hybrid learning technique (ANN-SNN Conversion, followed by STDB with trainable weights, threshold and leak) to train deep SNNs.

III. HYBRID SPIKE ENCODING

We propose a hybrid encoding scheme to convert the real-valued pixel intensities of input images into SNN inputs over the total number of timesteps dictated by the desired inference accuracy. As is typical, input images fed to the ANN are normalized to zero mean and unit standard deviation. In our proposed coding technique, we feed the analog pixel value in the input layer of the SNN in the 1^{st} timestep. Next, we convert the real-valued pixels into a spike train starting from the 2^{nd} timestep representing the same information. Considering a gray image with pixel intensity values in the range $[I_{min}, I_{max}]$, each input neuron encodes the temporal information of its' corresponding pixel value in a single spike time in the range $[2, T]$ where T is the total number of timesteps. The firing time of the i^{th} input neuron, T_i, is computed based on the i^{th} pixel intensity value, I_i, as follows

$$T_i = \lfloor T + \left(\frac{2 - T}{I_{max} - I_{min}} \right) \cdot (I_i - I_{min}) \rceil \quad (4)$$

where $\lfloor . \rceil$ represents the nearest integer function. Eq. 4 is represented as the point-slope form of the linear relationship shown in Fig. 2(b) and $\lfloor . \rceil$ is applied because T_i should be integral. Note that Eq. 4 also implies that the spike train starts from the 2^{nd} timestep 2. The encoded value of the i^{th} neuron in the input layer is thus expressed as
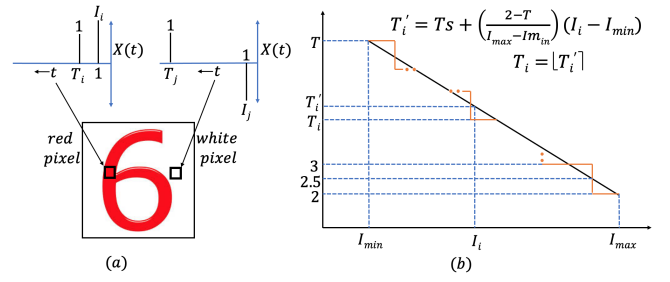

Fig. 2. (a) Hybrid coded input to the SNN (b) Mapping between the pixel intensity of images and the firing time of individual neurons where $\lfloor . \rceil$ denotes the nearest integer function

$$X_i(t) = \begin{cases} I_i, & \text{if } t = 1 \\ 1, & \text{else if } t = T_i \\ 0, & \text{otherwise} \end{cases} \quad (5)$$

which is further illustrated in Fig. 2(b). Brighter image pixels have higher intensities, and hence, lower T_i. Neurons at the subsequent layers fire as soon as they reach their threshold, and both the voltage potential and time to reach the threshold in the output layer determines the network decision. The analog pixel value in the 1^{st} time step influences the membrane potential of the output neurons, while the firing times of the input neurons based on the pixel intensities are responsible for the spike times of the output neurons.

Notably, this hybrid encoding scheme captures both the intensity and temporal nature of the input neurons, does not need any preprocessing steps like applying Gabor filters that are commonly used in SNNs trained with spike time dependent plasticity (STDP) [44], [45]. Moreover, our proposed encoding technique is compatible with event-driven cameras which capture actual pixel value first, and subsequently emit spikes based on the changes in pixel intensity [46]. Lastly, our proposal ensures that there is a single input spike per pixel, and hence, the obtained spike train is sparser than that observed in rate/direct coded techniques.

IV. PROPOSED TRAINING SCHEME

We employ a modified version of the LIF model illustrated in Section II to train energy-efficient SNNs. In our proposed training framework, neurons in all the hidden convolutional and fully-connected layers (except the output layer) spike at most once over all the timesteps. During inference, once a neuron emits a spike, it is shut off, and does not participate in the remaining LIF computations. However, during training, the neurons in the hidden layers follow the model illustrated in Eq. (6)-(8) which shows that even though each neuron can fire at most once, it still needs to perform computations following the LIF model. This ensures that the error gradients are still non-zero following the spike time and enables our proposed training framework to avoid the dead neuron problem where learning does not happen in the absence of a spike.

$$\boldsymbol{U}_l^t = \lambda_l \boldsymbol{U}_l^{t-1} + W_l \boldsymbol{O}_{l-1}^t - V_l \cdot (\boldsymbol{z}_l^{t-1} > 0) \quad (6)$$

$$\boldsymbol{z}_l^t = \frac{\boldsymbol{U}_l^t}{V_l} - 1 \quad (7)$$

$$\boldsymbol{O}_l^t = \begin{cases} 1, & \text{if } \boldsymbol{z}_l^t > 0 \text{ and } \boldsymbol{z}_l^{t_i} \leq 0 \ \forall t_i \in [1, t) \\ 0, & \text{otherwise} \end{cases} \quad (8)$$

Note that $\boldsymbol{U}_l^t$, $\boldsymbol{O}_{l-1}$, and W_l are vectors containing the membrane potential of the neurons of layer l at timestep

t, spike signals from layer $(l-1)$, and the weight matrix connecting the layer l and $(l-1)$. Also note that $(\boldsymbol{z}_l^{t-1} > 0)$ in Eq. 6 denotes a Boolean vector of size equal to the number of neurons in layer l. The leak and threshold voltage for all the neurons in layer l are represented by λ_l and V_l respectively. In our training framework, both these parameters (same for the neurons in a particular layer) are trained with backpropagation along with the weights to optimize both accuracy and latency.

The neurons in the output layer accumulate the incoming inputs without any leakage as shown in Eq. 9. However, unlike previous works [24], [30], the output neurons in our proposed framework emit spikes following a model shown in Eq. 10 where $\boldsymbol{T}_l$ denote the vector containing the spike times of the output neurons and T is the total number of timesteps.

$$\boldsymbol{U}_l^t = \boldsymbol{U}_l^{t-1} + W_l \boldsymbol{O}_{l-1}^t \tag{9}$$

$$\boldsymbol{T}_l = \begin{cases} T, & \text{if } \boldsymbol{U}_l^T < V_l \\ \boldsymbol{t} \ s.t. \ \boldsymbol{U}_l^t \geq V_l \ \& \ \boldsymbol{U}_l^{t-1} < V_l, & \text{otherwise} \end{cases} \tag{10}$$

The output layer only triggers an output spike if there was no spike in the earlier timesteps and the corresponding membrane potential crosses the threshold. Also, an output neuron is forced to fire at the last timestep if it was unable to emit a spike in any of the timesteps. This ensures that all the neurons in the output layer have a valid T_l which can be included in the loss function.

Let us now we derive the expressions to compute the gradients of the trainable parameters of all the layers. We perform the spatial and temporal credit assignment by unrolling the network in temporal axis and employing backpropagation through time (BPTT) [30].

Output Layer: The loss function is defined on both $\boldsymbol{U}_l^T$ and $\boldsymbol{T}_l$ to correctly capture both the direct and temporal information presented at the input layer. Therefore, we employ two softmax functions of the i^{th} output neuron shown in Eq. 11, where N denotes the total number of classes, U_i^T and t_i represent the accumulated membrane potential after the final timestep and the firing time of the i^{th} neuron respectively. .

$$\tilde{U}_i = \frac{e^{U_i^T}}{\sum_{j=1}^{N} e^{U_j^T}}, \quad \tilde{t}_i = \frac{e^{-t_i}}{\sum_{j=1}^{N} e^{-t_j}} \tag{11}$$

The resulting hybrid cross entropy loss ($\mathcal{L}$) and its' gradient with respect to the accumulated membrane potential vector ($\frac{\partial \mathcal{L}}{\partial \boldsymbol{U}_l^T}$) are thus defined as

$$\mathcal{L} = -\sum_{i=1}^{N} y_i log(\tilde{U}_i \tilde{t}_i), \quad \frac{\partial \mathcal{L}}{\partial \boldsymbol{U}_l^T} = \tilde{\boldsymbol{U}}_l^T - \boldsymbol{y} \tag{12}$$

where $\tilde{\boldsymbol{U}}_l^T$ is the vector containing the softmax values $\tilde{U}_i$, and $\boldsymbol{y}$ is the one-hot encoded vector of the correct class. Similarly, the gradient with respect to the firing time vector ($\frac{\partial \mathcal{L}}{\partial \boldsymbol{T}_l}$) is ($\tilde{\boldsymbol{T}}_l - \boldsymbol{y}$). Now, we compute the weight update as

$$W_l = W_l - \eta \Delta W_l \tag{13}$$

$$\Delta W_l = \sum_t \frac{\partial \mathcal{L}}{\partial W_l} = \sum_t \frac{\partial \mathcal{L}}{\partial \boldsymbol{U}_l^t} \frac{\partial \boldsymbol{U}_l^t}{\partial W_l}$$

$$= \frac{\partial \mathcal{L}}{\partial \boldsymbol{U}_l^T} \sum_t \frac{\partial \boldsymbol{U}_l^t}{\partial W_l} = (\tilde{\boldsymbol{U}}_l^T - \boldsymbol{y}) \sum_t \boldsymbol{O}_{l-1}^t \tag{14}$$

where η is the learning rate (LR). In order to evaluate the threshold update at the output layer, we rewrite Eq. (10) as

$$\boldsymbol{T}_l = \sum_{t=1}^{T-1} (t \mathcal{H}(a) \mathcal{H}(b)) + T \mathcal{H}(c) \tag{15}$$

where $\mathcal{H}$ denotes the Heaviside step function, $\boldsymbol{a} = \boldsymbol{U}_l^t - \boldsymbol{V}_l$, $\boldsymbol{b} = \boldsymbol{V}_l - \boldsymbol{U}_l^{t-1}$, and $\boldsymbol{c} = \boldsymbol{V}_l - \boldsymbol{U}_l^T$. Note that $\boldsymbol{V}_l$ represents a vector of repeated elements of the threshold voltage of the output layer. The derivative ($\frac{\partial \boldsymbol{T}_l}{\partial V_l}$) can then be represented as

$$\frac{\partial \boldsymbol{T}_l}{\partial V_l} = \sum_{t=1}^{T-1} t(\mathcal{H}(\boldsymbol{a})\delta(\boldsymbol{b}) - \mathcal{H}(\boldsymbol{b})\delta(\boldsymbol{a})) + T\delta(\boldsymbol{c}) \tag{16}$$

where δ represents the Dirac-delta function. Since the delta function is zero almost everywhere, it will not allow the gradient of $\boldsymbol{T}_l$ to change and train V_l. Hence, we approximate Eq. (16) as

$$\sum_{t=1}^{T-1} t(\mathcal{H}(\boldsymbol{a})(|\boldsymbol{b}|<\boldsymbol{\beta}) - \mathcal{H}(\boldsymbol{b})(|\boldsymbol{a}|<\boldsymbol{\beta})) + T(|\boldsymbol{c}|<\boldsymbol{\beta}) \tag{17}$$

where $\boldsymbol{\beta}$ is a vector of size equal to the number of output neurons, consisting of the repeated elements of a training hyperparameter that controls the gradient of $\boldsymbol{T}_l$. Note that $(|\boldsymbol{a}|<\boldsymbol{\beta})$, $(|\boldsymbol{b}|<\boldsymbol{\beta})$, and $(|\boldsymbol{c}|<\boldsymbol{\beta})$ are all Boolean vectors of the same size as $\boldsymbol{\beta}$. We then compute the threshold update as

$$V_l = V_l - \eta \Delta V_l, \quad \Delta V_l = \frac{\partial \mathcal{L}}{\partial V_l} = \frac{\partial \mathcal{L}}{\partial \boldsymbol{T}_l} \frac{\partial \boldsymbol{T}_l}{\partial V_l} \tag{18}$$

Hidden Layers: The weight update of the l^{th} hidden layer is calculated from Eq. (6)-(8) as

$$\Delta W_l = \sum_t \frac{\partial \mathcal{L}}{\partial W_l} = \sum_t \frac{\partial \mathcal{L}}{\partial \boldsymbol{z}_l^t} \frac{\partial \boldsymbol{z}_l^t}{\partial \boldsymbol{O}_l^t} \frac{\partial \boldsymbol{O}_l^t}{\partial \boldsymbol{U}_l^t} \frac{\partial \boldsymbol{U}_l^t}{\partial W_l}$$

$$= \sum_t \frac{\partial \mathcal{L}}{\partial \boldsymbol{z}_l^t} \frac{\partial \boldsymbol{z}_l^t}{\partial \boldsymbol{O}_l^t} \frac{1}{V_l} \boldsymbol{O}_{l-1}^t \tag{19}$$

$\frac{d\boldsymbol{z}_l^t}{d\boldsymbol{O}_l^t}$ is the non-differentiable gradient which can be approximated with the surrogate gradient proposed in [42].

$$\frac{\partial \boldsymbol{z}_l^t}{\partial \boldsymbol{O}_l^t} = \gamma \cdot max(0, 1 - |\boldsymbol{z}_l^t|) \tag{20}$$

where γ is a hyperparameter denoting the maximum value of the gradient. The threshold update is then computed as

$$\Delta V_l = \sum_t \frac{\partial \mathcal{L}}{\partial V_l} = \sum_t \frac{\partial \mathcal{L}}{\partial \boldsymbol{O}_l^t} \frac{\partial \boldsymbol{O}_l^t}{\partial \boldsymbol{z}_l^t} \frac{\partial \boldsymbol{z}_l^t}{\partial V_l}$$

$$= \sum_t \frac{\partial \mathcal{L}}{\partial \boldsymbol{O}_l^t} \frac{\partial \boldsymbol{O}_l^t}{\partial \boldsymbol{z}_l^t} \left(\frac{-V_l \cdot (\boldsymbol{z}_l^{t-1} > 0) - \boldsymbol{U}_l^t}{(V_l)^2} \right) \tag{21}$$

Given that the threshold is same for all neurons in a particular layer, it may seem redundant to train both the weights and threshold together. However, our experimental evaluation detailed in Section VI shows that the number of timesteps required to obtain the state-of-the-art classification accuracy decreases with this joint optimization. We hypothesize that this is because the optimizer is able to reach an improved local minimum when both parameters are tunable. Finally, the leak update is computed as

$$\lambda_l = \lambda_l - \eta \Delta \lambda_l \qquad (22)$$

$$\Delta \lambda_l = \sum_t \frac{\partial \mathcal{L}}{\partial \lambda_l} = \sum_t \frac{\partial \mathcal{L}}{\partial \boldsymbol{O}_l^t} \frac{\partial \boldsymbol{O}_l^t}{\partial \boldsymbol{z}_l^t} \frac{\partial \boldsymbol{z}_l^t}{\partial \boldsymbol{U}_l^t} \frac{\partial \boldsymbol{U}_l^t}{\partial \lambda_l}$$

$$= \sum_t \frac{\partial \mathcal{L}}{\partial \boldsymbol{O}_l^t} \frac{\partial \boldsymbol{O}_l^t}{\partial \boldsymbol{z}_l^t} \frac{1}{V_l} \boldsymbol{U}_l^{(t-1)} \qquad (23)$$

V. EXPERIMENTS

This section first describes how we evaluate the efficacy of our proposed encoding and training framework and then presents the inference accuracy on CIFAR-10 and CIFAR-100 datasets with various VGG model variants.

A. Experimental Setup

1) ANN Training for Initialization: To train our ANNs, we used the standard data-augmented input set for each model. For ANN training with various VGG models, we imposed a number of constraints that leads to near lossless SNN conversion [15]. In particular, our models are trained without the bias term because it complicates parameter space exploration which increases conversion difficulty and tends to increase conversion loss. The absence of bias term implies that Batch Normalization [47] cannot be used as a regularizer during the training process. Instead, we use Dropout [48] as the regularizer for both ANN and SNN training. Also, our pooling operations use average pooling because for binary spike based activation layers, max pooling incurs significant information loss. We performed the ANN training for 200 epochs with an initial LR of 0.01 that decays by a factor of 0.1 after 120, 160, and 180 epochs.

2) ANN-SNN Conversion and SNN Training: Previous works [15], [30] set the layer threshold of the first hidden layer by computing the maximum input to a neuron over all its neurons across all T timesteps for a set of input images [15]. The thresholds of the subsequent layers are sequentially computed in a similar manner taking the maximum across all neurons and timesteps. However, in our proposed framework, the threshold for each layer is computed sequentially as the 99.7 percentile (instead of the maximum) of the neuron input distribution at each layer, which improves the SNN classification accuracy [24]. During threshold computation, the leak in the hidden layers is set to unity and the analog pixel values of an image are directly applied to the input layer [24]. We considered only 512 input images to limit conversion time and used a threshold scaling factor of 0.4 for SNN training and inference, following the recommendations in [30].

Initialized with these layer thresholds and the trained ANN weights, we performed our proposed SNN training with the hybrid input encoding scheme for 150 epochs for CIFAR-10 and CIFAR-100, respectively, where we jointly optimize the weights, the membrane leak, and the firing thresholds of each layer as described in Section IV. We set $\gamma = 0.3$ [42], $\beta = 0.2$, and used a starting LR of 10^{-4} which decays by a factor of 0.1 every 10 epochs.

B. Classification Accuracy & Latency

We evaluated the performance of these networks on multiple VGG architectures, namely VGG-6, VGG-9 and VGG-11 for CIFAR-10 and VGG-16 for CIFAR-100 datasets respectively. Column-2 in Table I shows the ANN accuracy; column-3 shows the accuracy after ANN-SNN conversion with 200 timesteps. Note that we need 200 timesteps to evaluate the thresholds of the SNN for VGG architectures without any

TABLE I
MODEL PERFORMANCES WITH SINGLE-SPIKE HYBRID ENCODED SNN TRAINING ON CIFAR-10 AND CIFAR-100 AFTER A) ANN TRAINING, B) ANN-TO-SNN CONVERSION AND C) SNN TRAINING.

Architecture	a. ANN (%) accuracy	b. Accuracy (%) with ANN-SNN conversion for $T = 200$	c. Accuracy (%) after proposed SNN training for T=5
Dataset : CIFAR-10			
VGG-6	90.22	89.98	88.89
VGG-11	91.02	91.77	90.66
VGG-16	93.24	93.16	91.41
Dataset : CIFAR-100			
VGG-16	71.02	70.38	66.46

TABLE II
PERFORMANCE COMPARISON OF THE PROPOSED SINGLE SPIKE HYBRID ENCODED SNN WITH STATE-OF-THE-ART DEEP SNNs ON CIFAR-10 AND CIFAR-100. TTFS DENOTES TIME-TO-FIRST-SPIKE CODING.

Authors	Training type	Input encoding	Architecture	Accuracy (%)	Time steps
Dataset : CIFAR-10					
Sengupta et al. (2019) [15]	ANN-SNN conversion	Rate	VGG-16	91.55	2500
Wu et al. (2019) [27]	Surrogate gradient	Direct	5 CONV, 2 linear	90.53	12
Rathi et al. (2020) [30]	Conversion+ STDB training	Rate	VGG-16	91.13 92.02	100 200
Garg et al. (2019) [25]	Conversion+ STDB training	DCT	VGG-9	89.94	48
Kim et al. (2018) [21]	ANN-SNN conversion	Phase	VGG-16	91.2	1500
Park et al. (2019) [22]	ANN-SNN conversion	Burst	VGG-16	91.4	1125
Park et al. (2020) [18]	STDB training	TTFS	VGG-16	91.4	680
Kim at. al. (2020) [28]	Surrogate gradient	Rate	VGG-9	90.5	25
Rathi at. al. (2020) [24]	Conversion+ STDB training	Direct	VGG-16	92.70 **93.10**	5 10
This work	Conversion+ STDB training	Hybrid	VGG-16	91.41	5
Dataset : CIFAR-100					
Lu et al. (2020) [32]	ANN-SNN conversion	Direct	VGG-16	63.20	62
Garg et al. (2020) [25]	Conversion+ STDB training	DCT	VGG-11	68.3	48
Park et al. (2019) [22]	ANN-SNN conversion	Burst	VGG-16	68.77	3100
Park et al. (2020) [18]	STDB training	TTFS	VGG-16	68.8	680
Kim at. al. (2020) [28]	Surrogate gradient	Rate	VGG-9	66.6	50
Rathi et al. (2020) [24]	Conversion+ STDB training	Direct	VGG-16	**69.67**	5
This work	Conversion+ STDB training	Hybrid	VGG-16	66.46	5

significant loss in accuracy. Column-4 in Table I shows the accuracy when we perform our proposed training with our hybrid input encoding discussed in Section III. The performance of the SNNs trained via our proposed framework is compared with the current state-of-the-art SNNs with various encoding and training techniques in Table II. Our proposal requires only 5 timesteps for both SNN training and inference to obtain the SOTA test accuracy and hence, representing 5-300× improvement in inference latency compared to other rate/temporally coded spiking networks. Note that the direct encoding in the first time step is crucial for SNN convergence, and temporal coding solely leads to a test accuracy of ∼10% and ∼1% on CIFAR-10 and CIFAR-100 respectively, for all the network architectures.

VI. IMPROVEMENT IN ENERGY-EFFICIENCY

A. Reduction in Spiking Activity

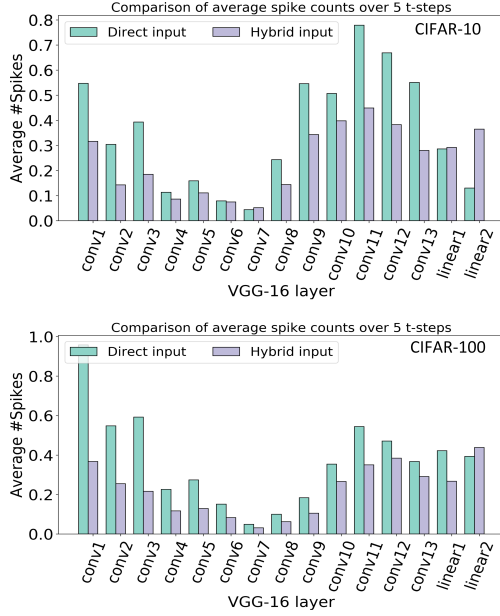To model energy consumption, we assume a generated SNN spike consumes a fixed amount of energy [12]. Based on this

Fig. 3. Comparison of average spiking activity per layer for VGG-16 on CIFAR-10 and CIFAR-100 with both direct and hybrid input encoding.
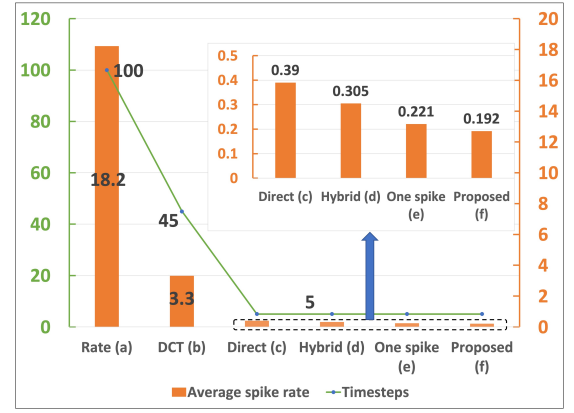

Fig. 4. Effect of Poisson rate encoding, DCT encoding, direct encoding, and single-spike hybrid input encoding on the average spike rate and latency for VGG-16 architecture on CIFAR-10 dataset.

SOTA test accuracy of all these SNNs. Both (d) and (e) result in lower average spiking activity compared to all the SOTA SNNs, with at most the same number of timesteps. Finally, (f) generates even lower number of average spikes ($2\times, 17.2\times$, and $94.8\times$ less compared to direct, DCT, and rate coding) with the lowest inference latency reported till date for deep SNN architectures [24], and no significant reduction in the test accuracy. The improvement stems from both the hybrid input encoding which reduces spiking activity in the initial few layers and our single-spike constraint which reduces the average spike rate throughout the network, particularly in the later layers. It becomes increasingly difficult for the membrane potential of the convolutional layers deep into the network to increase sufficient to emit a spike, due to the fact that the neurons in the earlier layers cannot fire multiple times and we need only 5 timesteps for classification.

TABLE III
CONVOLUTIONAL AND FULLY-CONNECTED LAYER FLOPs FOR ANN AND SNN MODELS

Model		Number of FLOPs	
	Notation	Convolutional layer l	Fully-connected layer l
ANN	F_{ANN}^l	$(k^l)^2 \times H_o^l \times W_o^l \times C_o^l \times C_i^l$	$f_i^l \times f_o^l$
SNN	F_{SNN}^l	$(k^l)^2 \times H_o^l \times W_o^l \times C_o^l \times C_i^l \times \zeta^l$	$f_i^l \times f_o^l \times \zeta^l$

B. Reduction in FLOPs and Compute Energy

Let us assume a convolutional layer l having weight tensor $\mathbf{W}^l \in \mathbb{R}^{k^l \times k^l \times C_i^l \times C_o^l}$ that operates on an input activation tensor $\mathbf{I}^l \in \mathbb{R}^{H_i^l \times W_i^l \times C_i^l}$, where H_i^l, W_i^l, C_i^l and C_o^l are the input tensor height, width, number of channels, and filters, respectively. k^l represents both filter height and width. We now quantify the energy consumed to produce the corresponding output activation tensor $\mathbf{O}^l \in \mathbb{R}^{H_o^l \times W_o^l \times C_o^l}$ for an ANN and SNN, respectively. Our model can be extended to fully-connected layers with f_i^l and f_o^l as the number of input and output features respectively. In particular, for an ANN, the total number of FLOPS for layer l, denoted F_{ANN}^l, is shown in row 1 of Table III [49], [50]. The formula can be easily adjusted for an SNN in which the number of FLOPs at layer l is a function of the average spiking activity at the layer (ζ^l) denoted as F_{SNN}^l in Table III. Thus, as the activation output gets sparser, the compute energy decreases.

For ANNs, FLOPs primary consist of multiply accumulate (MAC) operations of the convolutional and linear layers. On

---

assumption, earlier works [15], [30] have adopted the average spiking activity (also known as average spike count) of an SNN layer l, denoted ζ^l, as a measure of compute-energy of the model. In particular, ζ^l is computed as the ratio of the total spike count in T steps over all the neurons of the layer l to the total number of neurons in that layer. Thus lower the spiking activity, the better the energy efficiency.

Fig. 3 shows the average number of spikes for each layer with our proposed single-spike hybrid encoding and direct encoding scheme on VGG-16 when evaluated for 1500 samples from CIFAR-10 testset for VGG-16 architecture. Let the average be denoted by ζ^l which is computed by summing all the spikes in a layer over 100 timesteps and dividing by the number of neurons in that layer. For example, the average spike count of the 11^{th} convolutional layer of the direct encoded SNN is 0.78, which implies that over a 5 timestep period each neuron in that layer spikes 0.78 times on average over all input samples. As we can see, the spiking activity for almost all the layers reduces significantly with our proposed encoding technique.

To compare our proposed work with the SOTA SNNs, we perform hybrid training (ANN-SNN conversion, along with STDB) on spiking networks with (a) IF neurons with Poisson rate encoding [30], (b) IF neurons with DCT-based input encoding [25], and (c) LIF neurons with direct encoding [24]. We employ trainable leak and threshold in these SNNs for fair comparison. We also evaluate the impact of the hybrid spatio-temporal encoding with the modified loss function and the single-spike constraint individually on the average spike rate and latency under similar accuracy and conditions (trainable threshold and leak). In particular, we train three additional spiking networks: (d) SNN with LIF neuron and proposed hybrid-encoding, (e) SNN with LIF neuron and direct encoding with the single-spike constraint over all the layers, and (f) single-spike hybrid encoded SNN with LIF neuron. All the six networks achieve test accuracies between 90-93% for VGG-16 on CIFAR-10. Fig. 4 shows the average spiking rate and the number of timesteps required to obtain the

the contrary, for SNNs, except the first and last layer, the FLOPs are limited to accumulates (ACs) as the spikes are binary and thus simply indicate which weights need to be accumulated at the post-synaptic neurons. For the first layer, we need to use MAC units as we consume analog input[1] (at timestep one). Hence, the compute energy for an ANN (E_{ANN}) and an iso-architecture SNN model (E_{SNN}) can be written as

$$E_{ANN} = (\sum_{l=1}^{L} F^l_{SNN}) \cdot E_{MAC} \tag{24}$$

$$E_{SNN} = (F^1_{ANN}) \cdot E_{MAC} + (\sum_{l=2}^{L} F^l_{SNN}) \cdot E_{AC} \tag{25}$$

where L is the total number of layers. Note that E_{MAC} and E_{AC} are the energy consumption for a MAC and AC operation respectively. As shown in Table IV, E_{AC} is $\sim 32\times$ lower than E_{MAC} [51] in 45 nm CMOS technology. This number may vary for different technologies, but generally, in most technologies, an AC operation is significantly cheaper than a MAC operation.

Fig. 5 illustrates the energy consumption and FLOPs for ANN and SNN models of VGG-16 while classifying the CIFAR datasets, where the energy is normalized to that of an equivalent ANN. The number of FLOPs for SNNs obtained by our proposed training framework is smaller than that for an ANN with similar number of parameters. Moreover, because the ACs consume significantly less energy than MACs, as shown in Table IV, SNNs are significantly more energy efficient. In particular, for CIFAR-10 our proposed SNN consumes $\sim 70\times$ less compute energy than a comparable iso-architecture ANN with similar parameters and $\sim 1.2\times$ less compute energy than a comparable SNN with direct encoding technique and trainable threshold/leak [24] parameters. For CIFAR-100 with hybrid encoding and our single-spike constraint, the energy-efficiency can reach up to $\sim 125\times$ and $\sim 1.8\times$, respectively, compared to ANN and direct-coded SNN models [24] having similar parameters and architecture. Note that we did not consider the memory access energy in our evaluation because it is dependent on the underlying system architecture. Although SNNs incur significant data movement because the membrane potentials need to be fetched at every timestep, there have been many proposals to reduce the memory cost by data buffering [52], computing in non-volatile crossbar memory arrays [53], and data reuse with energy-efficient dataflows [9]. All these techniques can be applied to the SNNs obtained by our proposed training framework to address the memory cost.

VII. CONCLUSIONS

SNNs that operate with discrete spiking events can potentially unlock the energy wall in deep learning for edge applications. Towards this end, we presented a training framework that leads to low latency, energy-efficient spiking networks with high activation sparsity. We initialize the parameters of our proposed SNN taken from a trained ANN, to speed-up the training with spike-based backpropagation. The image pixels are applied directly as input to the network during the first

[1]Note that for the hybrid coded data input we need to perform MAC at the first layer at $t = 1$, and AC operation during remaining timesteps at that layer. For the direct coded input, only MAC during the 1^{st} timestep is sufficient, as neither the inputs nor the weights change during remaining timesteps (i.e. $5 \geq t \geq 2$).
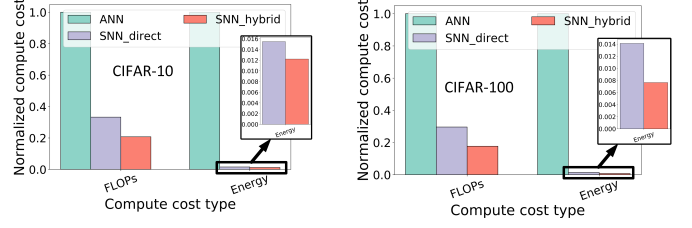

Fig. 5. Comparison of normalized compute cost on CIFAR-10 and CIFAR-100 for VGG-16 of ANN and SNN with direct and hybrid input encoding.

TABLE IV
ESTIMATED ENERGY COSTS FOR MAC AND AC OPERATIONS IN 45 nm CMOS PROCESS AT 0.9 V [51]

Serial No.	Operation	Energy (pJ)
1.	32-bit multiplication int	3.1
2.	32-bit addition int	0.1
3.	32-bit MAC	3.2 (#1 + #2)
4.	32-bit AC	0.1 (#2)

timestep, while they are converted to a sparse spike train with firing times proportional to the pixel intensities in subsequent timesteps. We also employ a modified version of the LIF model for the hidden and output layers of the SNN, in which all the neurons fire at most once per image. Both of these lead to high activation sparsity in the input, convolutional, and dense layers of the network. Moreover, we employ a hybrid cross entropy loss function to account for the spatio-temporal encoding in the input layer and train the network weights, firing threshold, and membrane leak via spike-based backpropagation to optimize both accuracy and latency. The high sparsity combined with low inference latency reduces the compute energy by ~ 70-$130\times$ and ~ 1.2-$1.8\times$ compared to an equivalent ANN and a direct encoded SNN respectively with similar accuracy. SNNs obtained by our proposed framework achieves similar accuracy as other state-of-the-art rate or temporally coded SNN models with 5-$300\times$ fewer timesteps. Future works include power and performance evaluation of our energy-efficient models on neurmorphic chips such as Loihi [54] and exploration of neuromorphic datasets [55] to leverage the temporal learning ability of our training framework.

VIII. ACKNOWLEDGEMENTS

REFERENCES

[1] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, pp. 436–44, 05 2015.

[2] l. Deng, J. Li, J.-T. Huang, K. Yao, D. Yu, F. Seide, M. Seltzer, G. Zweig, X. He, J. Williams, Y. Gong, and A. Acero, "Recent advances in deep learning for speech research at microsoft," in *Acoustics, Speech, and Signal Processing, 1988. ICASSP-88., 1988 International Conference on*, 10 2013, pp. 8604–8608.

[3] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in Neural Information Processing Systems 25*, F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, Eds. Curran Associates, Inc., 2012, pp. 1097–1105.

[4] C. Chen, A. Seff, A. Kornhauser, and J. Xiao, "DeepDriving: Learning affordance for direct perception in autonomous driving," in *2015 IEEE International Conference on Computer Vision (ICCV)*, vol. 1, no. 1, 2015, pp. 2722–2730.

[5] A. Rezvantalab, H. Safigholi, and S. Karimijeshni, "Dermatologist level dermoscopy skin cancer classification using different deep learning convolutional neural networks algorithms," *arXiv preprint arXiv:1810.10348*, 2018.

[6] S. Han, J. Pool, J. Tran, and W. J. Dally, "Learning both weights and connections for efficient neural networks," in *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1*. MIT Press, 2015, p. 1135–1143.

[7] S. Han, H. Mao, and W. J. Dally, "Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding," *arXiv preprint arXiv:1510.00149*, 2015.

[8] S. Lym, D. Lee, M. O'Connor, N. Chatterjee, and M. Erez, "DeLTA: GPU performance model for deep learning applications with in-depth memory system traffic analysis," *2019 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, Mar 2019.

[9] Y.-H. Chen, J. Emer, and V. Sze, "Eyeriss: A spatial architecture for energy-efficient dataflow for convolutional neural networks," in *ACM SIGARCH Computer Architecture News*, vol. 44, 06 2016.

[10] G. Indiveri and T. Horiuchi, "Frontiers in neuromorphic engineering," *Frontiers in Neuroscience*, vol. 5, 2011.

[11] M. Pfeiffer and T. Pfeil, "Deep learning with spiking neurons: Opportunities and challenges," *Frontiers in Neuroscience*, vol. 12, p. 774, 2018.

[12] Y. Cao, Y. Chen, and D. Khosla, "Spiking deep convolutional neural networks for energy-efficient object recognition," *International Journal of Computer Vision*, vol. 113, pp. 54–66, 05 2015.

[13] A. Sengupta, A. Banerjee, and K. Roy, "Hybrid spintronic-CMOS spiking neural network with on-chip learning: Devices, circuits, and systems," *Phys. Rev. Applied*, vol. 6, Dec 2016.

[14] P. U. Diehl, G. Zarrella, A. Cassidy, B. U. Pedroni, and E. Neftci, "Conversion of artificial recurrent neural networks to spiking neural networks for low-power neuromorphic hardware," in *2016 IEEE International Conference on Rebooting Computing (ICRC)*. IEEE, 2016, pp. 1–8.

[15] A. Sengupta, Y. Ye, R. Wang, C. Liu, and K. Roy, "Going deeper in spiking neural networks: VGG and residual architectures," *Frontiers in Neuroscience*, vol. 13, p. 95, 2019.

[16] I. M. Comsa, K. Potempa, L. Versari, T. Fischbacher, A. Gesmundo, and J. Alakuijala, "Temporal coding in spiking neural networks with alpha synaptic function," in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, vol. 1, no. 1, 2020, pp. 8529–8533.

[17] S. Zhou, X. LI, Y. Chen, S. T. Chandrasekaran, and A. Sanyal, "Temporal-coded deep spiking neural network with easy training and robust performance," *arXiv preprint arXiv:1909.10837*, 2020.

[18] S. Park, S. Kim, B. Na, and S. Yoon, "T2FSNN: Deep spiking neural networks with time-to-first-spike coding," *arXiv preprint arXiv:2003.11741*, 2020.

[19] M. Zhang, J. Wang, B. Amornpaisannon, Z. Zhang, V. Miriyala, A. Belatreche, H. Qu, J. Wu, Y. Chua, T. E. Carlson, and H. Li, "Rectified linear postsynaptic potential function for backpropagation in deep spiking neural networks," 2020.

[20] S. R. Kheradpisheh and T. Masquelier, "Temporal backpropagation for spiking neural networks with one spike per neuron," *International Journal of Neural Systems*, vol. 30, no. 06, May 2020.

[21] J. Kim, H. Kim, S. Huh, J. Lee, and K. Choi, "Deep neural networks with weighted spikes," *Neurocomputing*, vol. 311, pp. 373–386, 2018.

[22] S. Park, S. Kim, H. Choe, and S. Yoon, "Fast and efficient information transmission with burst spikes in deep spiking neural networks," in *2019 56th ACM/IEEE Design Automation Conference (DAC)*, vol. 1, no. 1, 2019, pp. 1–6.

[23] D. Almomani, M. Alauthman, M. Alweshah, O. Dorgham, and F. Albalas, "A comparative study on spiking neural network encoding schema: implemented with cloud computing," *Cluster Computing*, vol. 22, 06 2019.

[24] N. Rathi and K. Roy, "DIET-SNN: Direct input encoding with leakage and threshold optimization in deep spiking neural networks," *arXiv preprint arXiv:2008.03658*, 2020.

[25] I. Garg, S. S. Chowdhury, and K. Roy, "DCT-SNN: Using DCT to distribute spatial information over time for learning low-latency spiking neural networks," *arXiv preprint arXiv:2010.01795*, 2020.

[26] J. H. Lee, T. Delbruck, and M. Pfeiffer, "Training deep spiking neural networks using backpropagation," *Frontiers in Neuroscience*, vol. 10, 2016.

[27] Y. Wu, L. Deng, G. Li, J. Zhu, Y. Xie, and L. Shi, "Direct training for spiking neural networks: Faster, larger, better," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 2019, pp. 1311–1318.

[28] Y. Kim and P. Panda, "Revisiting batch normalization for training low-latency deep spiking neural networks from scratch," *arXiv preprint arXiv:2010.01729*, 2020.

[29] A. Tavanaei, M. Ghodrati, S. R. Kheradpisheh, T. Masquelier, and A. Maida, "Deep learning in spiking neural networks," *Neural Networks*, vol. 111, p. 47–63, Mar 2019.

[30] N. Rathi, G. Srinivasan, P. Panda, and K. Roy, "Enabling deep spiking neural networks with hybrid conversion and spike timing dependent backpropagation," *arXiv preprint arXiv:2005.01807*, 2020.

[31] S. Kundu, G. Datta, M. Pedram, and P. A. Beerel, "Spike-thrift: Towards energy-efficient deep spiking neural networks by limiting spiking activity via attention-guided compression," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, January 2021, pp. 3953–3962.

[32] S. Lu and A. Sengupta, "Exploring the connection between binary and spiking neural networks," *arXiv preprint arXiv:2002.10064*, 2020.

[33] C. Lee, S. S. Sarwar, P. Panda, G. Srinivasan, and K. Roy, "Enabling spike-based backpropagation for training deep neural network architectures," *Frontiers in Neuroscience*, vol. 14, p. 119, 2020.

[34] S. S. Chowdhury, C. Lee, and K. Roy, "Towards understanding the effect of leak in spiking neural networks," *arXiv preprint arXiv:2006.08761*, 2020.

[35] E. Ledinauskas, J. Ruseckas, A. Juršėnas, and G. Buračas, "Training deep spiking neural networks," *arXiv preprint arXiv:2006.04436*, 2020.

[36] B. Rueckauer, I.-A. Lungu, Y. Hu, M. Pfeiffer, and S.-C. Liu, "Conversion of continuous-valued deep networks to efficient event-driven networks for image classification," *Frontiers in Neuroscience*, vol. 11, p. 682, 2017.

[37] P. U. Diehl, D. Neil, J. Binas, M. Cook, S. Liu, and M. Pfeiffer, "Fast-classifying, high-accuracy spiking deep networks through weight and threshold balancing," in *2015 International Joint Conference on Neural Networks (IJCNN)*, vol. 1, no. 1, 2015, pp. 1–8.

[38] Y. Hu, H. Tang, and G. Pan, "Spiking deep residual network," *arXiv preprint arXiv:1805.01352*, 2018.

[39] P. O'Connor, D. Neil, S.-C. Liu, T. Delbruck, and M. Pfeiffer, "Real-time classification and sensor fusion with a spiking deep belief network," *Frontiers in neuroscience*, vol. 7, p. 178, 2013.

[40] C. Lee, P. Panda, G. Srinivasan, and K. Roy, "Training deep spiking convolutional neural networks with STDP-based unsupervised pre-training followed by supervised fine-tuning," *Frontiers in Neuroscience*, vol. 12, 2018.

[41] P. Panda and K. Roy, "Unsupervised regenerative learning of hierarchical features in spiking deep networks for object recognition," *arXiv preprint arXiv:1602.01510*, 2016.

[42] G. Bellec, D. Salaj, A. Subramoney, R. Legenstein, and W. Maass, "Long short-term memory and learning-to-learn in networks of spiking neurons," *arXiv preprint arXiv:1803.09574*, 2018.

[43] E. O. Neftci, H. Mostafa, and F. Zenke, "Surrogate gradient learning in spiking neural networks: Bringing the power of gradient-based optimization to spiking neural networks," *IEEE Signal Processing Magazine*, vol. 36, no. 6, pp. 51–63, 2019.

[44] S. R. Kheradpisheh, M. Ganjtabesh, S. J. Thorpe, and T. Masquelier, "STDP-based spiking deep convolutional neural networks for object recognition," *Neural Networks*, vol. 99, p. 56–67, Mar 2018. [Online]. Available: http://dx.doi.org/10.1016/j.neunet.2017.12.005

[45] M. Mozafari, S. R. Kheradpisheh, T. Masquelier, A. Nowzari-Dalini, and M. Ganjtabesh, "First-spike-based visual categorization using reward-modulated STDP," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 29, no. 12, pp. 6178–6190, 2018.

[46] G. Gallego, T. Delbruck, G. M. Orchard, C. Bartolozzi, B. Taba, A. Censi, S. Leutenegger, A. Davison, J. Conradt, K. Daniilidis, and D. Scaramuzza, "Event-based vision: A survey," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 1, no. 1, pp. 1–1, 2020.

[47] N. Bjorck, C. P. Gomes, B. Selman, and K. Q. Weinberger, "Understanding batch normalization," in *Advances in Neural Information Processing Systems*, 2018, pp. 7694–7705.

[48] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *Journal of Machine Learning Research*, vol. 15, pp. 1929–1958, 06 2014.

[49] S. Kundu, M. Nazemi, M. Pedram, K. M. Chugg, and P. A. Beerel, "Pre-defined sparsity for low-complexity convolutional neural networks," *IEEE Transactions on Computers*, vol. 69, no. 7, pp. 1045–1058, 2020.

[50] S. Kundu, S. Prakash, H. Akrami, P. A. Beerel, and K. M. Chugg, "pSConv: A pre-defined sparse kernel based convolution for deep CNNs," in *2019 57th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 2019, pp. 100–107.

[51] M. Horowitz, "1.1 Computing's energy problem (and what we can do about it)," in *2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*. IEEE, 2014, pp. 10–14.

[52] Y. Shen, M. Ferdman, and P. Milder, "Escher: A cnn accelerator with flexible buffering to minimize off-chip transfer," in *2017 IEEE 25th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, vol. 1, no. 1, 2017, pp. 93–100.

[53] B. Chen, F. Cai, J. Zhou, W. Ma, P. Sheridan, and W. D. Lu, "Efficient in-memory computing architecture based on crossbar arrays," in *2015 IEEE International Electron Devices Meeting (IEDM)*, vol. 1, no. 1, 2015, pp. 1–4.

[54] M. Davies, N. Srinivasa, T. H. Lin, G. Chinya, Y. Cao, S. H. Choday, G. Dimou, P. Joshi, N. Imam, S. Jain, Y. Liao, C. K. Lin, A. Lines, R. Liu, D. Mathaikutty, S. McCoy, A. Paul, J. Tse, G. Venkataramanan, Y. H. Weng, A. Wild, Y. Yang, and H. Wang, "Loihi: A neuromorphic manycore processor with on-chip learning," *IEEE Micro*, vol. 38, no. 1, pp. 82–99, 2018.

[55] S. Miao, G. Chen, X. Ning, Y. Zi, K. Ren, Z. Bing, and A. Knoll, "Neuromorphic vision datasets for pedestrian detection, action recognition, and fall detection," *Frontiers in Neurorobotics*, vol. 13, p. 38, 2019.