



Enabling Spike-Based Backpropagation for Training Deep Neural Network Architectures

Chankyu Lee^{*†}, Syed Shakib Sarwar^{*†}, Priyadarshini Panda, Gopalakrishnan Srinivasan and Kaushik Roy

Nanoelectronics Research Laboratory, School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN, United States

OPEN ACCESS

Edited by:

Hesham Mostafa,
University of California, San Diego,
United States

Reviewed by:

Bodo Rückauer,
ETH Zürich, Switzerland
Eric Hunsberger,
University of Waterloo, Canada
David Kappel,
Dresden University of Technology,
Germany

*Correspondence:

Chankyu Lee
lee2216@purdue.edu
Syed Shakib Sarwar
sarwar@purdue.edu

[†]These authors have contributed
equally to this work

Specialty section:

This article was submitted to
Neuromorphic Engineering,
a section of the journal
Frontiers in Neuroscience

Received: 11 September 2019

Accepted: 30 January 2020

Published: 28 February 2020

Citation:

Lee C, Sarwar SS, Panda P,
Srinivasan G and Roy K (2020)
Enabling Spike-Based
Backpropagation for Training Deep
Neural Network Architectures.
Front. Neurosci. 14:119.
doi: 10.3389/fnins.2020.00119

Spiking Neural Networks (SNNs) have recently emerged as a prominent neural computing paradigm. However, the typical shallow SNN architectures have limited capacity for expressing complex representations while training deep SNNs using input spikes has not been successful so far. Diverse methods have been proposed to get around this issue such as converting off-the-shelf trained deep Artificial Neural Networks (ANNs) to SNNs. However, the ANN-SNN conversion scheme fails to capture the temporal dynamics of a spiking system. On the other hand, it is still a difficult problem to directly train deep SNNs using input spike events due to the discontinuous, non-differentiable nature of the spike generation function. To overcome this problem, we propose an approximate derivative method that accounts for the leaky behavior of LIF neurons. This method enables training deep convolutional SNNs directly (with input spike events) using spike-based backpropagation. Our experiments show the effectiveness of the proposed spike-based learning on deep networks (VGG and Residual architectures) by achieving the best classification accuracies in MNIST, SVHN, and CIFAR-10 datasets compared to other SNNs trained with a spike-based learning. Moreover, we analyze sparse event-based computations to demonstrate the efficacy of the proposed SNN training method for inference operation in the spiking domain.

Keywords: spiking neural network, convolutional neural network, spike-based learning rule, gradient descent backpropagation, leaky integrate and fire neuron

1. INTRODUCTION

Over the last few years, deep learning has made tremendous progress and has become a prevalent tool for performing various cognitive tasks such as object detection, speech recognition, and reasoning. Various deep learning techniques (LeCun et al., 1998; Srivastava et al., 2014; Ioffe and Szegedy, 2015) enable the effective optimization of deep ANNs by constructing multiple levels of feature hierarchies and show remarkable results, which occasionally outperform human-level performance (Krizhevsky et al., 2012; He et al., 2016; Silver et al., 2016). To harness the deep learning capabilities in ubiquitous environments, it is necessary to deploy deep learning not only on large-scale computers, but also on edge devices (e.g., phone, tablet, smartwatch, robot, etc.). However, the ever-growing complexity of deep neural networks together with the explosion in the amount of data to be processed, place significant energy demands on current computing platforms.

Spiking Neural Network (SNN) is one of the leading candidates for overcoming the constraints of neural computing and to efficiently harness the machine learning algorithm in real-life (or mobile) applications. The concepts of SNN, which is often regarded as the 3rd generation neural network (Maass, 1997), are inspired by the biological neuronal mechanisms (Hodgkin and Huxley, 1952; Dayan and Abbott, 2001; Izhikevich, 2003; Brette and Gerstner, 2005) that can efficiently process discrete spatio-temporal events (spikes). The Leaky Integrate and Fire (LIF) neuron is the simple first-order phenomenological spiking neuron model, which can be characterized by the internal state, called membrane potential. The membrane potential integrates the inputs over time and generates an output spike whenever it overcomes the neuronal firing threshold. Recently, specialized hardwares (Merolla et al., 2014; Ankit et al., 2017; Davies et al., 2018) have been developed to exploit this event-based, asynchronous signaling/processing scheme. They are promising for achieving ultra-low power intelligent processing of streaming spatiotemporal data, and especially in deep hierarchical networks, as it has been observed in SNN models that the number of spikes, and thus the amount of computation, decreases significantly at deeper layers (Rueckauer et al., 2017; Sengupta et al., 2019).

We can divide SNNs into two broad classes: (a) converted SNNs and (b) SNNs derived by direct spike-based training. The former one is SNNs converted from the trained ANN for the efficient event-based inference (ANN-SNN conversion) (Cao et al., 2015; Hunsberger and Eliasmith, 2015; Diehl et al., 2016; Rueckauer et al., 2017; Sengupta et al., 2019). The main advantage is that it uses state-of-the-art (SOTA), optimization-based, ANN training methods and therefore achieves SOTA classification performance. For instance, the specialized SNN hardwares [such as SpiNNaker (Furber et al., 2013), IBM TrueNorth (Merolla et al., 2014) have exhibited greatly improved power efficiency as well as the state-of-the-art performance for the inference. The signals used in such training are real-valued and are naturally viewable as representing spike rate (frequency). The problem is that reliably estimating frequencies requires non-trivial passage of time. On the other hand, SNNs derived by direct spike-based training also involves some hurdles. Direct spike-based training methods can be divided into two classes: (i) non-optimization-based, mostly unsupervised, approaches involving only signals local to the synapse, e.g., the times of pre- and post-synaptic spikes, as in Spike-Timing-Dependent-Plasticity (STDP); and (ii) optimization-based, mostly supervised, approaches involving a global objective, e.g., loss, function. STDP-trained two-layer network (consisting of 6,400 output neurons) has been shown to achieve 95% classification accuracy on MNIST dataset (Diehl and Cook, 2015). However, a shallow network structure limits the expressive power (Brader et al., 2007; Zhao et al., 2015; Srinivasan et al., 2018a,b) and may not scale well to larger problem sizes. While efficient feature extraction has been demonstrated using layer-wise STDP learning in deep convolutional SNNs (Kheradpisheh et al., 2016; Lee et al., 2019b), ANN models trained with standard backpropagation (BP) (Rumelhart et al., 1985) still achieve significantly better classification performance. These considerations have inspired the search for spike-based

versions of BP, which requires finding a differentiable surrogate of the spiking unit's activation function. Bohte et al. (2002) and Lee et al. (2016) defined such a surrogate in terms of a unit's membrane potential. Besides, Panda and Roy (2016) applies BP-based supervised training for the classifier after layer-by-layer autoencoder-based training of the feature extractor. By combining layer-wise STDP-based unsupervised and supervised spike-based BP, Lee et al. (2018) showed improved robustness, generalization ability, and faster convergence. In this paper, we take these prior works forward to effectively train very deep SNNs using end-to-end spike-based BP learning.

The main contributions of our work are as follows. First, we develop a spike-based supervised gradient descent BP algorithm that employs an approximate (pseudo) derivative for LIF neuronal function. In addition, we leverage the key idea of the successful deep ANN models such as LeNet5 (LeCun et al., 1998), VGG (Simonyan and Zisserman, 2014), and ResNet (He et al., 2016) for efficiently constructing deep convolutional SNN architectures. We also adapt the dropout (Srivastava et al., 2014) technique to better regularize deep SNN training. Next, we demonstrate the effectiveness of our methodology for visual recognition tasks on standard character and object datasets (MNIST, SVHN, CIFAR-10) and a neuromorphic dataset (N-MNIST). To the best of our knowledge, this work achieves the best classification accuracy in MNIST, SVHN, and CIFAR-10 datasets among other spike-based learning methodologies. Last, we expand our efforts to quantify and analyze the advantages of a spike-based BP algorithm compared to ANN-SNN conversion techniques in terms of inference time and energy consumption.

The rest of the paper is organized as follows. In section 2.1, we provide the background on fundamental components and architectures of deep convolutional SNNs. In section 2.2.1, we detail the spike-based gradient descent BP learning algorithm. In section 2.2.2, we describe our spiking version of the dropout technique. In section 3.1–3.2, we describe the experiments and report the simulation results, which validate the efficacy of spike-based BP training for MNIST, SVHN, CIFAR-10, and N-MNIST datasets. In section 4.1, we discuss the proposed algorithm in comparison to relevant works. In section 4.2–4.4, we analyze the spike activity, inference speedup and complexity reduction of directly trained SNNs and ANN-SNN converted networks. Finally, we summarize and conclude the paper in section 5.

2. MATERIALS AND METHODS

2.1. The Components and Architecture of Spiking Neural Network

2.1.1. Spiking Neural Network Components

Leaky-Integrate-and-Fire (LIF) neurons (Dayan and Abbott, 2001) and plastic synapses are fundamental and biologically plausible computational elements for emulating the dynamics of SNNs. The sub-threshold dynamics of a LIF spiking neuron can be formulated as

$$\tau_m \frac{dV_{mem}}{dt} = -V_{mem} + I(t) \quad (1)$$

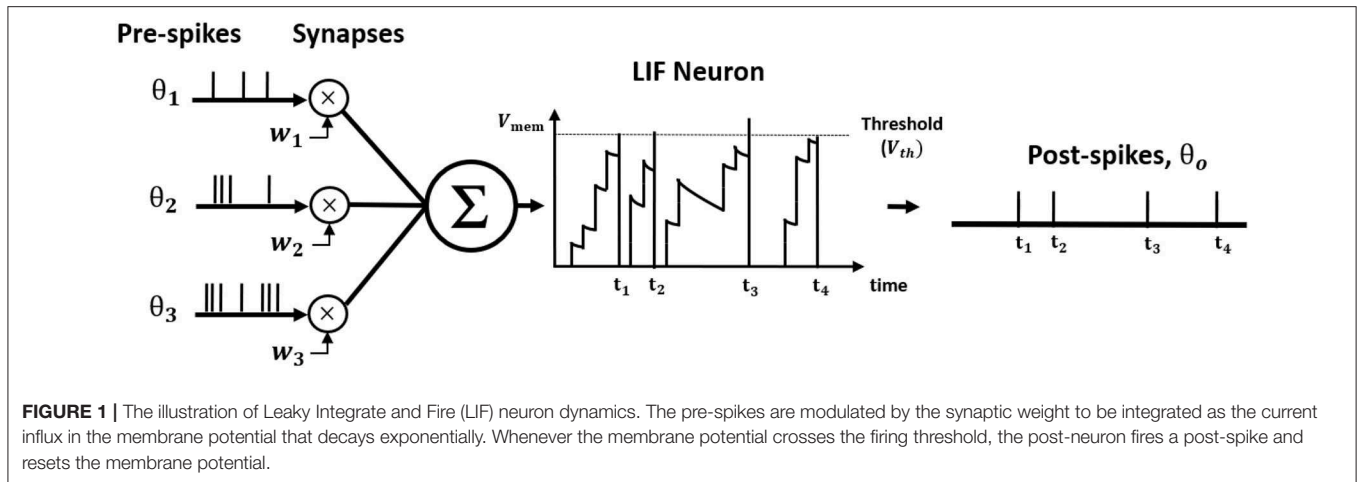


FIGURE 1 | The illustration of Leaky Integrate and Fire (LIF) neuron dynamics. The pre-spikes are modulated by the synaptic weight to be integrated as the current influx in the membrane potential that decays exponentially. Whenever the membrane potential crosses the firing threshold, the post-neuron fires a post-spike and resets the membrane potential.

where V_{mem} is the post-neuronal membrane potential and τ_m is the time constant for membrane potential decay. The input current, $I(t)$, is defined as the weighted summation of pre-spikes at each time step as given below.

$$I(t) = \sum_{i=1}^{n^l} (w_i \sum_k \theta_i(t - t_k)) \quad (2)$$

where n^l indicates the number of pre-synaptic weights, w_i is the synaptic weight connecting i th pre-neuron to post-neuron. $\theta_i(t - t_k)$ is a spike event from i th pre-neuron at time t_k , which can be formulated as a Kronecker delta function as follows,

$$\theta(t - t_k) = \begin{cases} 1, & \text{if } t = t_k \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

where t_k is the time instant that k th spike occurred. **Figure 1** illustrates LIF neuronal dynamics. The impact of each pre-spike, $\theta_i(t - t_k)$, is modulated by the corresponding synaptic weight (w_i) to generate a current influx to the post-neuron. Note, the units do not have bias term. The input current is integrated into the post-neuronal membrane potential (V_{mem}) that leaks exponentially over time with time constant (τ_m). When the membrane potential exceeds a threshold (V_{th}), the neuron generates a spike and resets its membrane potential to initial value. The **Table 1** lists the annotations used in Equations (1–27).

2.1.2. Deep Convolutional Spiking Neural Network

2.1.2.1. Building blocks

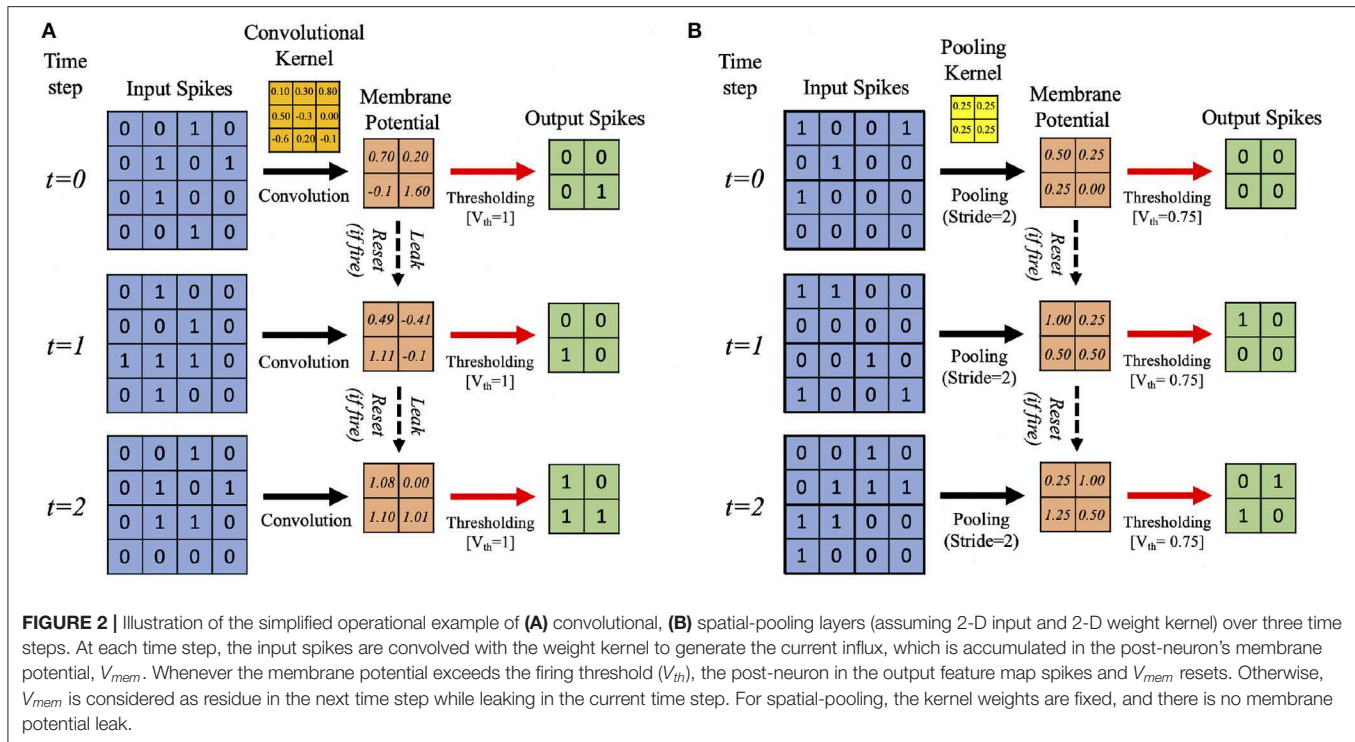
In this work, we develop a training methodology for convolutional SNN models that consist of an input layer followed by intermediate hidden layers and a final classification layer. In the input layer, the pixel images are encoded as Poisson-distributed spike trains where the probability of spike generation is proportional to the pixel intensity. The hidden layers consist of multiple convolutional (C) and spatial-pooling (P) layers, which are often arranged in an alternating manner. These convolutional (C) and spatial-pooling (P) layers represent

TABLE 1 | List of notations.

Notations	Meaning
θ	Spike event
x	The sum of pre-spike events over time
w	Synaptic weight
V_{mem}	Membrane potential
V_{th}	Neuronal firing threshold
I	Input current at each time step
net	Total current influx over time
a	Activation of spiking neuron
E	Loss function
δ	Error gradient

the intermediate stages of feature extractor. The spikes from the feature extractor are combined to generate a one-dimensional vector input for the fully-connected (FC) layers to produce the final classification. The convolutional and fully-connected layers contain trainable parameters (i.e., synaptic weights) while the spatial-pooling layers are fixed *a priori*. Through the training procedure, weight kernels in the convolutional layers can encode the feature representations of the input patterns at multiple hierarchical levels. **Figure 2A** shows the simplified operational example of a convolutional layer consisting of LIF neurons over three time steps (assuming 2-D input and 2-D weight kernel). On each time step, each neuron convolves its input spikes with the weight kernel to compute its input current, which is integrated into its membrane potential, V_{mem} . If $V_{mem} > V_{th}$, the neuron spikes and its V_{mem} is set to 0. Otherwise, V_{mem} is considered as residue in the next time step while leaking in the current time step. **Figure 2B** shows the simplified operation of a pooling layer, which reduces the dimensionality from the previous convolutional layer while retaining spatial (topological) information.

The two major operations used for pooling are max and average. Both have been used for SNNs, e.g., max-pooling (Rueckauer et al., 2017) and average-pooling (Cao et al., 2015;

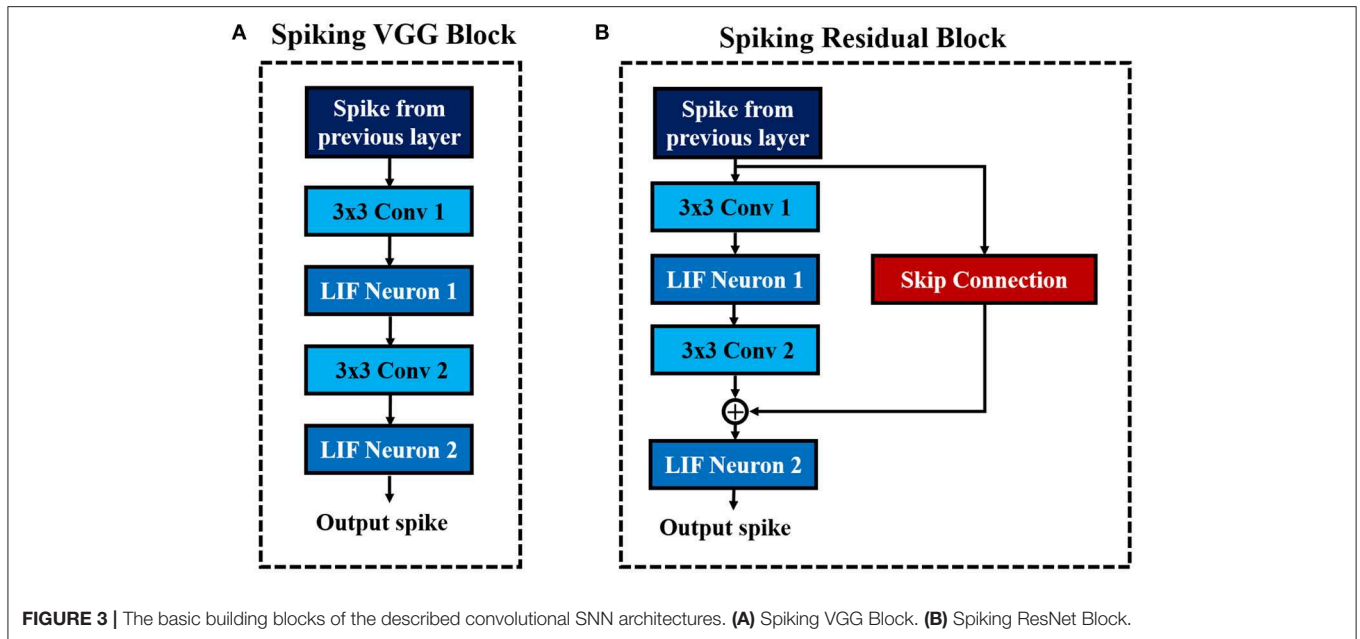


Diehl et al., 2015). We use average-pooling due to its simplicity. In the case of SNNs, an additional thresholding is used after averaging to generate output spikes. For instance, a fixed 2×2 kernel (each having a weight of 0.25) strides through a convolutional feature map without overlapping and fires an output spike at the corresponding location in the pooled feature map only if the sum of the weighted spikes of the 4 inputs within the kernel window exceeds a designated firing threshold (set to 0.75 in this work). Otherwise, the membrane potential remains as a residue in the next time step. **Figure 2B** shows an example spatial-pooling operation over three time steps (assuming 2-D input and 2-D weight kernel). The average-pooling threshold need to be carefully set so that spike propagation is not disrupted due to the pooling. If the threshold is too low, there will be too many spikes, which can cause loss of spatial location of the feature that was extracted from the previous layer. If the threshold is too high, there will not be enough spike propagation to the deeper layers.

2.1.2.2. Deep convolutional SNN architecture: VGG and residual SNNs

Deep networks are essential for recognizing intricate input patterns so that they can effectively learn hierarchical representations. To that effect, we investigate popular deep neural network architectures such as VGG (Simonyan and Zisserman, 2014) and ResNet (He et al., 2016) in order to build deep SNN architectures. VGG (Simonyan and Zisserman, 2014) was one of the first neural networks, which used the idea of using small (3×3) convolutional kernels uniformly throughout the network. Using small kernels enables effective stacking of convolutional layers while minimizing the number of parameters

in deep networks. In this work, we build deep convolutional SNNs (containing more than 5 trainable layers) using “Spiking VGG Blocks,” which contain stacks of convolutional layers using small (3×3) kernels. **Figure 3A** shows a “Spiking VGG block” containing two stacked convolutional layers, each followed by a LIF neuronal layer. The convolutional layer box contains the synaptic connectivity, and the LIF neuronal box contains the activation units. Next, ResNet (He et al., 2016) introduced the skip connections throughout the network that had considerable successes in enabling successful training of significantly deeper networks. In particular, ResNet addresses the degradation (of training accuracy) problem (He et al., 2016) that occurs while increasing the number of layers in the standard feedforward neural network. We employ the concept of skip connection to construct deep residual SNNs with 7–11 trainable layers. **Figure 3B** shows a “Spiking Residual Block” containing non-residual and residual paths. The non-residual path consists of two convolutional layers with an intermediate LIF neuronal layer. The residual path (skip connection) is composed of the identity mapping when the number of input and output feature maps are the same, and 1×1 convolutional kernels when the number of input and output feature maps differ. The outputs of both the non-residual and residual paths are integrated to the membrane potential in the last LIF neuronal activation layer (LIF Neuron 2 in **Figure 3B**) to generate output spikes from the “Spiking Residual Block.” Within the feature extractor, a “Spiking VGG Block” or “Spiking Residual Block” is often followed by an average-pooling layer. Note, in some “Spiking Residual Blocks,” the last convolutional and residual connections employ convolution with a stride of 2 to incorporate the functionality of the spatial-pooling layers. At the end of the feature extractor,



extracted features from the last average-pooling layer are fed to a fully-connected layer as a 1-D vector input for inference.

2.2. Supervised Training of Deep Spiking Neural Network

2.2.1. Spike-Based Gradient Descent Backpropagation Algorithm

The spike-based BP algorithm in SNN is adapted from standard BP (Rumelhart et al., 1985) in the ANN domain. In standard BP, the network parameters are iteratively updated in a direction to minimize the difference between the final outputs of the network and target labels. The standard BP algorithm achieves this goal by back-propagating the output error through the hidden layers using gradient descent. However, the major difference between ANNs and SNNs is the dynamics of neuronal output. An artificial neuron (such as *sigmoid*, *tanh*, or *ReLU*) communicates via continuous values whereas a spiking neuron generates binary spike outputs over time. In SNNs, spatiotemporal spike trains are fed to the network as inputs. Accordingly, the outputs of spiking neuron are spike events, which are also discrete over time. Hence, the standard BP algorithm is incompatible with training SNNs, as it can not back-propagate the gradient through a non-differentiable spike generation function. In this work, we formulate an approximate derivative for LIF neuron activation, making gradient descent possible. We derive a spike-based BP algorithm that is capable of learning spatiotemporal patterns in spike-trains. The spike-based BP can be divided into three phases, forward propagation, backward propagation and weight update, which we describe in the following sections.

2.2.1.1. Forward propagation

In forward propagation, spike trains representing input patterns are presented to the network for estimating the network outputs. To generate the spike inputs, the input pixel values are converted

to Poisson-distributed spike trains and delivered to the network. The input spikes are multiplied with synaptic weights to produce an input current that accumulates in the membrane potential of post neurons as in Equation (1). Whenever its membrane potential exceeds a neuronal firing threshold, the post-neuron generates an output spike and resets. Otherwise, the membrane potential decays exponentially over time. The neurons of every layer (excluding output layer) carry out this process successively based on the weighted spikes received from the preceding layer. Over time, the total weighted summation of the pre-spike trains (i.e., net) is described as follows,

$$net_j^l(t) = \sum_{i=1}^{n^{l-1}} (w_{ij}^{l-1} x_i^{l-1}(t)), \text{ where } x_i^{l-1}(t) = \sum_t \sum_k \theta_i^{l-1}(t - t_k) \quad (4)$$

where $net_j^l(t)$ represents the total current influx integrated to the membrane potential of j th post-neuron in layer l over the time t , n^{l-1} is the number of pre-neurons in layer $l-1$ and $x_i^{l-1}(t)$ denotes the sum of spike train ($t_k \leq t$) from i th pre-neuron over time t . The sum of post-spike trains ($t_k \leq t$) is represented by $a_j^l(t)$ for the j th post-neuron.

$$a_j^l(t) = \sum_t \sum_k \theta_j^l(t - t_k) \quad (5)$$

Clearly, the sum of post-spike train ($a^l(t)$) is equivalent to the sum of pre-spike train ($x^l(t)$) for the next layer. On the other hand, the neuronal firing threshold of the final classification layer is set to a very high value so that final output neurons do not spike. In the final layer, the weighted pre-spikes are accumulated in the membrane potential while decaying over time. At the last time step, the accumulated membrane potential is divided by the

number of total time steps (T) in order to quantify the output distribution (*output*) as presented by Equation (6).

$$output = \frac{V_{mem}^L(T)}{number\ of\ timesteps} \tag{6}$$

2.2.1.2. Backward propagation and weight update

Next, we describe the backward propagation for the proposed spike-based backpropagation algorithm. After the forward propagation, the loss function is measured as a difference between target labels and outputs predicted by the network. Then, the gradients of the loss function are estimated at the final layer. The gradients are propagated backward all the way down to the input layer through the hidden layers using recursive chain rule, as formulated in Equation (7). The following Equations (7–27) and **Figure 4** describe the detailed steps for obtaining the partial derivatives of (final) output error with respect to weight parameters.

The prediction error of each output neuron is evaluated by comparing the output distribution (*output*) with the desired target label (*label*) of the presented input spike trains, as shown in Equation (8). The corresponding loss function (*E* in Equation, 9) is defined as the sum of squared (final prediction) error over all output neurons. To calculate the $\frac{\partial E}{\partial a_{LIF}}$ and $\frac{\partial a_{LIF}}{\partial net}$ terms in Equation (7), we need a defined activation function and a method to differentiate the activation function of a LIF neuron.

$$\frac{\partial E}{\partial w^l} = \frac{\partial E}{\partial a_{LIF}} \frac{\partial a_{LIF}}{\partial net} \frac{\partial net}{\partial w^l} \tag{7}$$

$$Final\ output\ error,\ e_j = output_j - label_j \tag{8}$$

$$Loss\ function,\ E = \frac{1}{2} \sum_{j=1}^{n^l} e_j^2 \tag{9}$$

In SNN, the “activation function” indicates the relationship between the weighted summation of pre-spike inputs and post-neuronal outputs over time. In forward propagation, we have different types of neuronal activation for the final layer and hidden layers. Hence, the estimation of neuronal activations and their derivatives are different for the final layer and hidden layers. For the final layer, the value of *output* in Equation (6) is used as the neuronal activation (a_{LIF}) while considering the discontinuities at spike time instant as noise. Hence, $\frac{\partial E}{\partial output}$ is equal to the final output error, as calculated in Equation (10).

$$\frac{\partial E}{\partial output} = \frac{\partial}{\partial output} \frac{1}{2} (output - label)^2 = output - label = e \tag{10}$$

During back-propagating phase, we consider the leak statistics of membrane potential in the final layer neurons as noise. This allows us to approximate the accumulated membrane potential value for a given neuron as equivalent to the total input current (i.e., *net*) received by the neuron over the forward time duration (T) ($V_{mem,j}^L(T) \approx \sum_{i=1}^{n^l-1} (w_{ij}x_i(T)) = net_j^L(T)$). Therefore, the

derivative of post-neuronal activation with respect to *net* for final layer ($\frac{\partial output}{\partial net} \equiv \frac{\partial V_{mem}^L(T)/T}{\partial net} = \frac{\partial net^L(T)/T}{\partial net} = \frac{1}{T}$) is calculated as $\frac{1}{T}$ for the final layer.

For the hidden layers, we have post-spike trains as the neuronal outputs. The spike generation function is non-differentiable since it creates a discontinuity (because of step jump) at the time instance of firing. Hence, we introduce a pseudo derivative method for LIF neuronal activation ($a'_{LIF}(net)$) for the hidden layers, for back-propagating the output error via the chain rule. The purpose of deriving $a'_{LIF}(net)$ is to approximately estimate the $\frac{\partial a_{LIF}}{\partial net}$ term in Equation (7) for the hidden layers only. To obtain this pseudo derivative of LIF neuronal activation with respect to total input current (i.e., *net*), we make the following approximations. We first estimate the derivative of an “Integrate and Fire” (IF) neuron’s activation. Next, with the derivative of IF neuron’s activation, we estimate a leak correctional term to compensate for the leaky effect of membrane potential in LIF activation. Finally, we obtain an approximate derivative for LIF neuronal activation as a combination of two estimations (i.e., derivative for IF neuron and approximated leak compensation derivative). If a hidden neuron does not fire any spike, the derivative of corresponding neuronal activation is set to zero.

The spike generation function of IF neuron is a hard threshold function that generates the output signal as either +1 or 0. The IF neuron fires a post-spike whenever the input currents accumulated in membrane potential exceed the firing threshold (note, in case of IF neuron, there is no leak in the membrane potential). Hence, the membrane potential of a post-neuron at time instant *t* can be written as,

$$V_{mem}(t) \approx \sum_{i=1}^n (w_i x_i(t)) - V_{th} a_{IF}(t) \tag{11}$$

where *n* denotes the number of pre-neurons, $x_i(t)$ is the sum of spike events from *i*th pre-neuron over time *t* (defined in Equation, 4) and $a_{IF}(t)$ represents the sum of post-spike trains over time *t* (defined in Equation 5). In Equation (11), $\sum_{i=1}^n (w_i x_i(t))$ accounts for the integration behavior and $V_{th} a_{IF}(t)$ accounts for the fire/reset behavior of the membrane potential dynamics. If we assume V_{mem} as zero (using small signal approximation), the activation of IF neuron ($a_{IF}(t)$) can be formulated as the Equation (12). Then, by differentiating it with respect to *net* (in Equation, 13), the derivative of IF neuronal activation can be approximated as a linear function with slope of $\frac{1}{V_{th}}$ as the straight-through estimation (Bengio et al., 2013).

$$a_{IF}(t) \approx \frac{1}{V_{th}} \sum_{i=1}^n (w_i x_i(t)) = \frac{1}{V_{th}} net(t) \tag{12}$$

$$\frac{\partial a_{IF}}{\partial net} \approx \frac{1}{V_{th}} 1 = \frac{1}{V_{th}} \tag{13}$$

The spike generation function of both the IF and LIF neuron models are the same, namely the hard threshold function.

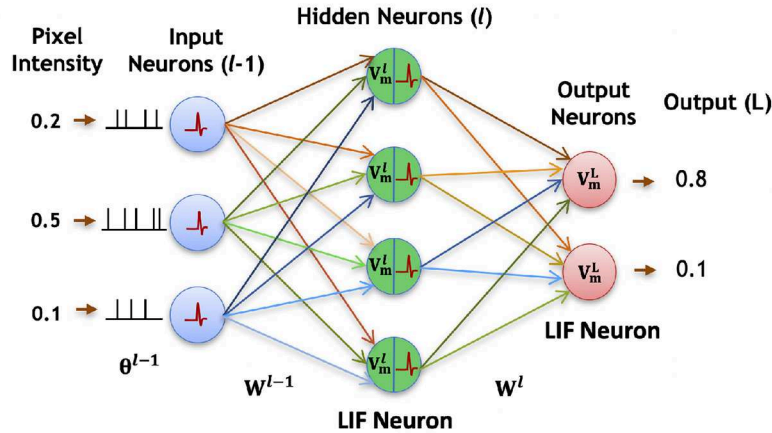
Forward Pass

• Hidden Layer

Total input current : $net^l(t) = \sum_{i=1}^{n^{l-1}} W_i^{l-1} X_i^{l-1}, X_i^{l-1} = \sum_t \sum_k \theta_i^{l-1}(t-t_k)$
 Activation of Neurons : $a_{LIF}^l(t) = \sum_t \sum_k \theta^l(t-t_k)$

• Final Layer

$net^L(t) = \sum_{i=1}^{n^L} W_i^L X_i^L, X_i^L = \sum_t \sum_k \theta_i^L(t-t_k)$
 $a_{LIF}^L(net^L) = output = \frac{1}{T} V_{mem}^L(t)$



Backward Pass

• Hidden Layer

Error Gradient : $\delta^l = ((W^l)^T \delta^L) \cdot a_{LIF}^l'(net^l)$
 $a_{LIF}^l'(net^l) = \frac{1}{V_{th}} (1 + \frac{1}{Y} f(t))$

• Final Layer

$\delta^L = (output - label) \cdot a_{LIF}^L'(net^L) = e \cdot \frac{1}{T}$
 Output Error (e) = output - label

FIGURE 4 | Illustration of the forward and backward propagation phase of the proposed spike-based BP algorithm in a multi-layer SNN comprised of LIF neurons. In the forward phase, the LIF neurons (in all layers) accumulate the weighted sum of the pre-spikes in the membrane potential, which decays exponentially over time. In addition, the LIF neurons in hidden layers generate post-spikes if the membrane potential exceeds a threshold and reset the membrane potential. However, the LIF neurons in the final layer, do not generate any spike, but rather accumulate the weighted sum of pre-spikes till the last time step to quantify the final outputs. Then, the final errors are evaluated by comparing the final outputs to the label data. In the backward phase, the final errors are propagated backward through the hidden layers using the chain rule to obtain the partial derivatives of final error with respect to weights. Finally, the synaptic weights are modified in a direction to reduce the final errors.

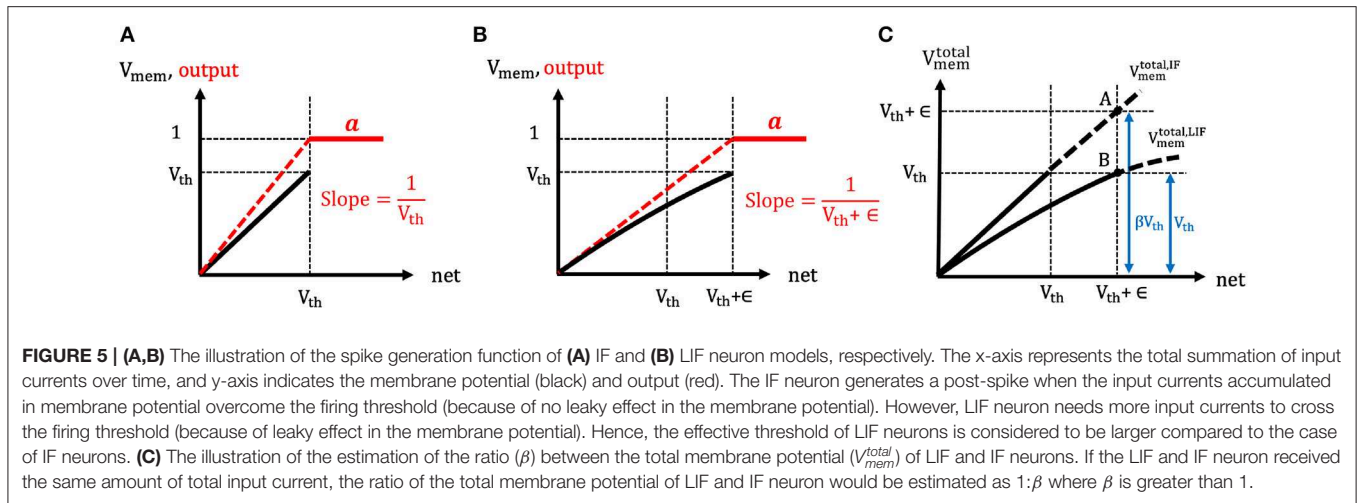
However, the effective neuronal thresholds are considered to be different for the two cases, as shown in **Figures 5A,B**. In the LIF neuron model, due to the leaky effect in the membrane potential, larger input current (as compared to IF neuron) needs to be accumulated in order to cross the neuronal threshold and generate a post-spike. Hence, the effective neuronal threshold becomes $V_{th} + \epsilon$ where ϵ is a positive value that reflects the leaky effect of membrane potential dynamics. Now, the derivative of LIF neuronal activation ($\frac{\partial a_{LIF}}{\partial net}$) can be approximated as a hard threshold function [similar to IF and Equation (13)] and written as $\frac{1}{V_{th} + \epsilon}$. Clearly, the output of a LIF neuron depends on the firing threshold and leaky characteristics (embodied in ϵ) of the membrane potential whereas the output of an IF neuron depends only on the firing threshold. Next, we explain the detailed steps to estimate the ϵ and in turn calculate the derivative of LIF neuronal activation ($\frac{\partial a_{LIF}}{\partial net}$).

To compute ϵ , the ratio (β) between the total membrane potential ($V_{mem}^{total}(t)$) of IF and LIF neurons is estimated at the end of forward propagation time (T) as shown in **Figure 5C**. Here, $V_{mem}^{total}(t)$ represents the hypothetical total membrane potential with accumulated input current without reset mechanism

until time step (t). Suppose both the IF and LIF neurons received the same amount of total input current (i.e., $net(T)$), the total membrane potential of LIF neuron is expected to be lower than the total membrane potential of IF neuron ($V_{mem}^{total,LIF}(T) : V_{mem}^{total,IF}(T) = 1 : \beta$ where $\beta > 1$). Hence, by comparing the total membrane potential values of IF and LIF neurons in **Figure 5C**, the relation of ϵ and β can be obtained as follows,

$$V_{th} + \epsilon = \beta V_{th} \tag{14}$$

where $V_{th} + \epsilon$ represents the total membrane potential of IF neuron (point A in **Figure 5C**) and V_{th} indicates the total membrane potential of LIF neuron (point B in **Figure 5C**) when both neurons received the same amount of net inputs. Based on this assumption, we now estimate the ratio (β) by using the relation of the spike output evolution ($\frac{\partial a(t)}{\partial t}$) and the total membrane potential evolution ($\frac{\partial V_{mem}^{total}(t)}{\partial t}$) over time as described in Equations (16–20). As mentioned previously, the total input current (i.e., $net(t)$) and total membrane potential ($V_{mem}^{total}(t)$) are estimated similar to that of IF neuron (because



of no leaky effect) so that Equation (15) can be derived from Equation (12). By differentiating Equation (15) with respect to time, we get the relation of the spike output evolution ($\frac{\partial a_{IF}(t)}{\partial t}$) and the membrane potential evolution ($\frac{\partial V_{mem}^{total,IF}(t)}{\partial t}$) over time for IF neuron as described in Equation (16).

$$a_{IF}(t) \approx \frac{1}{V_{th}} net(t) \approx \frac{1}{V_{th}} V_{mem}^{total,IF}(t) \quad (15)$$

$$\frac{\partial a_{IF}(t)}{\partial t} \approx \frac{1}{V_{th}} \frac{\partial V_{mem}^{total,IF}(t)}{\partial t} \quad (16)$$

Hence, in IF neuron case, the evolution of membrane potential over time ($\frac{\partial V_{mem}^{total,IF}(t)}{\partial t}$) can be represented by the multiplication of firing threshold (V_{th}) and the spike output evolution ($\frac{\partial a_{IF}(t)}{\partial t}$) in Equation (17). Note, the evolution of membrane potential over time ($\frac{\partial V_{mem}^{total,IF}(t)}{\partial t}$) indicates the integration component due to the average input current over time. We consider $a_{IF}(t)$ as homogeneous spike trains where spike firing rates are constant, so that the $\frac{\partial a_{IF}(t)}{\partial t}$ can be replaced with the post-neuronal firing rate ($rate(t)$). The homogeneous post-neuronal firing rate, $rate(t)$, can be represented by $\frac{a(t)}{t}$ where $a(t)$ is the number of post-spikes and t means the given forward time window. In LIF neuron case, however, the evolution of membrane potential ($\frac{\partial V_{mem}^{total,LIF}(t)}{\partial t}$) can be expressed as the combination of average input current (integration component) and leaky (exponential decay) effect as shown in Equation (18). To measure the leaky effect in Equation (18), we estimate the low-pass filtered output spikes ($t_k \leq t$) that leak over time using the function $V_{th}f(t)$ (depicted in Equation, 19), and differentiate it with respect to time at $t \rightarrow t_k^+$ (from the right-sided limit). The $V_{th}f(t)$, as a post-synaptic potential, contains the total membrane potential history over time. The time constant (τ_m) in Equation (19) determines the decay rate of post-synaptic potential. Essentially, the main idea is to approximately estimate the leaky effect by comparing the total

membrane potential and obtain the ratio (β) between both cases (i.e., IF and LIF neurons).

$$\frac{\partial V_{mem}^{total,IF}(t)}{\partial t} \approx V_{th} \frac{\partial a_{IF}(t)}{\partial t} \approx V_{th} rate(t) \quad (17)$$

$$\frac{\partial V_{mem}^{total,LIF}(t)}{\partial t} \approx V_{th} rate(t) + V_{th} \frac{\partial f(t)}{\partial t} \quad (18)$$

$$f(t) = \sum_k \exp\left(-\frac{t-t_k}{\tau_m}\right) \quad (19)$$

$$\frac{\partial a_{IF}(t)}{\partial t} \approx \frac{1}{V_{th}} \frac{\partial V_{mem}^{total,IF}(t)}{\partial t} = \beta \frac{1}{V_{th}} \frac{\partial V_{mem}^{total,LIF}(t)}{\partial t} \quad (20)$$

By solving the Equations (17, 18, 20), the inverse ratio ($\frac{1}{\beta}$) is derived as follows in Equation (21),

$$\frac{1}{\beta} = 1 + \frac{1}{rate(t)} \frac{\partial f(t)}{\partial t} \quad (21)$$

where the first term (unity) indicates the effect of average input currents (that is observed from the approximate derivative of IF neuron activation, namely the straight-through estimation) and the second term ($\frac{1}{rate(t)} \frac{\partial f(t)}{\partial t}$) represents the leaky (exponential decay) effect of LIF neuron for the forward propagation time window. Then, by using the relations of ϵ and β in Equation (14), the derivative of LIF neuronal activation can be obtained as $\frac{\partial a_{LIF}}{\partial net} = \frac{1}{V_{th} + \epsilon} = \frac{1}{\beta V_{th}}$. In this work, to avoid the vanishing gradient phenomena during the error back-propagation, the leaky effect term ($\frac{1}{rate(t)} \frac{\partial f(t)}{\partial t}$) is divided by the size of the forward propagation time window (T). Hence, the scaled time derivative of this function, $\frac{1}{\gamma} f'(t)$, is used as the leak correctional term where γ denotes the number of output spike events for a particular neuron over the total forward propagation time. As a result, we obtain an approximate derivative for LIF neuronal

activation (in hidden layers) as a combination of the straight-through estimation (i.e., approximate derivative of IF neuron activation) and the leak correctional term that compensates leaky effect in the membrane potential as described in Equation (22). Please note that, in our work, input and output spikes are not exponentially decaying, the leak only happens according to the mechanism of membrane potential. Moreover, $f(t)$ is not a part of the forward propagation phase, and rather it is only defined to approximately measure the leaky effect during the backward propagation phase by differentiating it with respect to time. The function $f(t)$ is a time-dependent function that simply integrates the output spikes ($t_k \leq t$) temporally, and the resultant sum is decayed over time. It is evident that $f(t)$ is continuous except where spikes occur and the activities jump up (Lee et al., 2016). Therefore, $f(t)$ is differentiable at $t \rightarrow t_k^+$ (from the right-sided limit). Note that, to capture the leaky effect (exponential decay), it is necessary to compute the derivative of $f(t)$ at the points in between the spiking activities, not at the time instant of spiking.

$$\begin{aligned} \frac{\partial a_{LIF}}{\partial net} &= \frac{1}{V_{th} + \epsilon} = \frac{1}{\beta V_{th}} \approx \frac{1}{V_{th}} \left(1 + \frac{1}{\gamma} f'(t)\right) \\ &= \frac{1}{V_{th}} \left(1 + \frac{1}{\gamma} \sum_k -\frac{1}{\tau_m} e^{-\frac{t-t_k}{\tau_m}}\right) \end{aligned} \quad (22)$$

In summary, the approximations applied to implement a spike-based BP algorithm in SNN are as follows:

- During the back-propagating phase, we consider the leaks in the membrane potential of final layer neurons as noises so that the accumulated membrane potential is approximated as equivalent to the total input current ($V_{mem}^L \approx net$). Therefore, the derivative of post-neuronal activation with respect to net ($\frac{\partial output}{\partial net}$) is calculated as $\frac{1}{T}$ for the final layer.
- For hidden layers, we first approximate the activation of an IF neuron as a linear function (i.e., straight-through estimation). Hence, we are able to estimate its derivative of IF neuron's activation (Bengio et al., 2013) with respect to total input current.
- To capture the leaky effect of a LIF neuron (in hidden layers), we estimate the scaled time derivative of the low-pass filtered output spikes that leak over time, using the function $f(t)$. This function is continuous except for the time points where spikes occur (Lee et al., 2016). Hence, it is differentiable in the sections between the spiking activities.
- We obtain an approximate derivative for LIF neuronal activation (in hidden layers) as a combination of two derivatives. The first one is the straight-through estimation (i.e., approximate derivative of IF neuron activation). The second one is the leak correctional term that compensates the leaky effect in the membrane potential of LIF neurons. The combination of straight-through estimation and the leak correctional term is expected to be less than 1.

Based on these approximations, we can train SNNs with direct spike inputs using a spike-based BP algorithm.

At the final layer, the error gradient, δ^L , represents the gradient of the output loss with respect to total input current (i.e., net) received by the post-neurons. It can be calculated by multiplying

the final output error (e) with the derivative of the corresponding post-neuronal activation ($\frac{\partial output}{\partial net^L}$) as shown in Equation (23). At any hidden layer, the local error gradient, δ^l , is recursively estimated by multiplying the back-propagated gradient from the following layer ($(w^l)^{Tr} * \delta^{l+1}$) with derivative of the neuronal activation, $a'_{LIF}(net^l)$, as presented in Equation (24). Note that element-wise multiplication is indicated by “.” while matrix multiplication is represented by “*” in the respective equations.

$$\delta^L = \frac{\partial E}{\partial output} \frac{\partial output}{\partial net^L} = e \frac{1}{T} = \frac{e}{T} \quad (23)$$

$$\delta^l = ((w^l)^{Tr} * \delta^{l+1}).a'_{LIF}(net^l) \quad (24)$$

The derivative of net with respect to weight is simply the total incoming spikes over time as derived in Equation (25). The derivative of the output loss with respect to the weights interconnecting the layers l and $l + 1$ (Δw^l in Equation, 26) is determined by multiplying the transposed error gradient at $l + 1$ (δ^{l+1}) with the input spikes from layer l . Finally, the calculated partial derivatives of loss function are used to update the respective weights using a learning rate (η_{BP}) as illustrated in Equation (27). As a result, iterative updating of the weights over mini-batches of input patterns leads the network state to a local minimum, thereby enabling the network to capture multiple-levels of internal representations of the data.

$$\frac{\partial net}{\partial w^l} = \frac{\partial}{\partial w^l} (w^l * x^l(t)) = x^l(t) \quad (25)$$

$$\Delta w^l = \frac{\partial E}{\partial w^l} = x^l(t) * (\delta^{l+1})^{Tr} \quad (26)$$

$$w^l_{updated} = w^l - \eta_{BP} \Delta w^l \quad (27)$$

2.2.2. Dropout in Spiking Neural Network

Dropout (Srivastava et al., 2014) is one of the popular regularization techniques while training deep ANNs. This technique randomly disconnects certain units with a given probability (p) to avoid units being overfitted and co-adapted too much to given training data. There are prior works (Kappel et al., 2015, 2018; Neftci et al., 2015) that investigated the biological insights on how synaptic stochasticity can provide dropout-like functional benefits in SNNs. In this work, we employ the concept of dropout technique in order to regularize deep SNNs effectively. Note, dropout technique is only applied during training and is not used when evaluating the performance of the network during inference. There is a subtle difference in the way dropout is applied in SNNs compared to ANNs. In ANNs, each epoch of training has several iterations of mini-batches. In each iteration, randomly selected units (with dropout ratio of p) are disconnected from the network while weighting by its posterior probability ($\frac{1}{1-p}$). However, in SNNs, each iteration has more than one forward propagation depending on the time length of

Algorithm 1: Forward propagation with dropout at each iteration in SNN

```

1: Input : Poisson-distributed input spike train (inputs),
   Dropout ratio (p), Total number of time steps (#timesteps),
   Membrane potential ( $V_{mem}$ ), Time constant of membrane
   potential ( $\tau_m$ ), Firing threshold ( $V_{th}$ )
2: Initialize  $SNN^l.V_{mem} \leftarrow 0 \forall l = 2, \dots, \#SNN.layer$ 
3: // Define the random subset of units (with a probability  $1-p$ )
   at each iteration
4: for  $l \leftarrow 1$  to  $\#SNN.layer - 1$  do
5:    $mask^l \leftarrow generate\_random\_subset(probability = 1 - p)$ 
6: for  $t \leftarrow 1$  to  $\#timesteps$  do
7:   // Set input of first layer equal to spike train of a mini-
   batch data
8:    $SNN^1.spike[t] \leftarrow inputs[t]$ ;
9:   for  $l \leftarrow 2$  to  $\#SNN.layer$  do
10:    // Integrate weighted sum of input spikes to
   membrane potential
11:     $SNN^l.V_{mem}[t] \leftarrow SNN^l.V_{mem}[t - 1] +$ 
 $SNN^{l-1}.forward(SNN^{l-1}.spike[t]) \cdot (mask^{l-1} / (1-p))$ ;
12:    // If  $V_{mem}$  is greater than  $V_{th}$ , post-neuron generate a
   spike
13:    if  $SNN^l.V_{mem}[t] > SNN^l.V_{th}$  then
14:      // Membrane potential resets if the corresponding
   neuron fires a spike
15:       $SNN^l.spike[t] \leftarrow 1$ 
16:       $SNN^l.V_{mem}[t] \leftarrow 0$ 
17:    else
18:      // Else, membrane potential decays over time
19:       $SNN^l.spike[t] \leftarrow 0$ 
20:       $SNN^l.V_{mem}[t] \leftarrow e^{-\frac{1}{\tau_m}} * SNN^l.V_{mem}[t]$ 

```

the spike train. We back-propagate the output error and modify the network parameters only at the last time step. For dropout to be effective in our training method, it has to be ensured that the set of connected units within an iteration of mini-batch data is not changed, such that the neural network is constituted by the same random subset of units during each forward propagation within a single iteration. On the other hand, if the units are randomly connected at each time-step, the effect of dropout will be averaged out over the entire forward propagation time within an iteration. Then, the dropout effect would fade-out once the output error is propagated backward and the parameters are updated at the last time step. Therefore, we need to keep the set of randomly connected units for the entire time window within an iteration. In the experiment, we use the SNN version of dropout technique with the probability (p) of omitting units equal to 0.2–0.25. Note that the activations are much sparser in SNN forward propagations compared to ANNs, hence the optimal p for SNNs needs to be less than a typical ANN dropout ratio ($p = 0.5$). The details of SNN forward propagation with dropout are specified in **Algorithm 1**.

TABLE 2 | Parameters used in the experiments.

Parameter	Value
Time Constant of Membrane Potential (τ_m)	100 time-steps
BP Training Time Duration	50–100 time-steps
Inference Time Duration	Same as training
Mini-batch Size	16–32
Spatial-pooling Non-overlapping Region/Stride	2×2, 2
Neuronal Firing Threshold	1 (hidden layer), ∞ (final layer)
Weight Initialization Constant (κ)	2 (non-residual network), 1 (residual network)
Learning rate (η_{BP})	0.002–0.003
Dropout Ratio (p)	0.2–0.25

3. RESULTS

3.1. Experimental Setup

The primary goal of our experiments is to demonstrate the effectiveness of the proposed spike-based BP training methodology in a variety of deep network architectures. We first describe our experimental setup and baselines. For the experiments, we developed a custom simulation framework using the Pytorch (Paszke et al., 2017) deep learning package for evaluating our proposed SNN training algorithm. Our deep convolutional SNNs are populated with biologically plausible LIF neurons (with neuronal firing thresholds of unity) in which a pair of pre- and post- neurons are interconnected by plastic synapses. At the beginning, the synaptic weights are initialized with Gaussian random distribution of zero-mean and standard deviation of $\sqrt{\frac{\kappa}{n^l}}$ (n^l : number of fan-in synapses) as introduced in He et al. (2015). Note, the initialization constant κ differs by the type of network architecture. For instance, we have used $\kappa = 2$ for non-residual network and $\kappa = 1$ for residual network. For training, the synaptic weights are trained with a mini-batch spike-based BP algorithm in an end-to-end manner, as explained in section 2.2.1. For static datasets, we train our network models for 150 epochs using mini-batch stochastic gradient descent BP that reduces its learning rate at 70, 100, and 125th training epochs. For the neuromorphic dataset, we use Adam (Kingma and Ba, 2014) learning method and reduce its learning rate at 40, 80, and 120th training epochs. Please, refer to **Table 2** for more implementation details. The datasets and network topologies used for benchmarking, the input spike generation scheme for event-based operation and determination of the number of time-steps required for training and inference are described in the following sub-sections.

3.1.1. Benchmarking Datasets

We demonstrate the efficacy of our proposed training methodology for deep convolutional SNNs on three standard vision datasets and one neuromorphic vision dataset, namely the MNIST (LeCun et al., 1998), SVHN (Netzer et al., 2011), CIFAR-10 (Krizhevsky and Hinton, 2009), and N-MNIST (Orchard et al., 2015). The MNIST dataset is composed of gray-scale

TABLE 3 | Benchmark datasets.

Dataset	Image	#Training samples	#Testing samples	#Category
MNIST	28 × 28, gray	60,000	10,000	10
SVHN	32 × 32, color	73,000	26,000	10
CIFAR-10	32 × 32, color	50,000	10,000	10
N-MNIST	34 × 34 × 2 ON and OFF spikes	60,000	10,000	10

(one-dimensional) images of handwritten digits whose sizes are 28 by 28. The SVHN and CIFAR-10 datasets are composed of color (three-dimensional) images whose sizes are 32 by 32. The N-MNIST dataset is a neuromorphic (spiking) dataset that is converted from static MNIST dataset using Dynamic Vision Sensor (DVS) (Lichtsteiner et al., 2008). The N-MNIST dataset contains two-dimensional images that include ON and OFF event stream data whose sizes are 34 by 34. The ON (OFF) event represents the increase (decrease) in pixel bright changes. The details of the benchmark datasets are listed in **Table 3**. For evaluation, we report the top-1 classification accuracy by classifying the test samples (training samples and test samples are mutually exclusive).

3.1.2. Network Topologies

We use various SNN architectures depending on the complexity of the benchmark datasets. For MNIST and N-MNIST datasets, we used a network consisting of two sets of alternating convolutional and spatial-pooling layers followed by two fully-connected layers. This network architecture is derived from LeNet5 model (LeCun et al., 1998). Note that **Table 4** summarizes the layer type, kernel size, the number of output feature maps, and stride of SNN model for MNIST dataset. The kernel size shown in the table is for 3-D convolution where the 1st dimension is for number of input feature-maps and 2nd–3rd dimensions are for convolutional kernels. For SVHN and CIFAR-10 datasets, we used deeper network models consisting of 7 to 11 trainable layers including convolutional, spatial-pooling and fully-connected layers. In particular, these networks consisting of beyond 5 trainable layers are constructed using small (3 × 3) convolutional kernels. We term the deep convolutional SNN architecture that includes 3 × 3 convolutional kernel (Simonyan and Zisserman, 2014) without residual connections as “VGG SNN” and with skip (residual) connections (He et al., 2016) as “Residual SNN.” In Residual SNNs, some convolutional layers convolve kernel with the stride of 2 in both x and y directions, to incorporate the functionality of spatial-pooling layers. Please, refer to **Tables 4, 5** that summarize the details of deep convolutional SNN architectures. In the results section, we will discuss the benefit of deep SNNs in terms of classification performance as well as inference speedup and energy efficiency.

3.1.3. ANN-SNN Conversion Scheme

As mentioned previously, off-the-shelf trained ANNs can be successfully converted to SNNs by replacing ANN (ReLU)

neurons with Integrate and Fire (IF) spiking neurons and adjusting the neuronal thresholds with respect to synaptic weights. In the literature, several methods have been proposed (Cao et al., 2015; Hunsberger and Eliasmith, 2015; Diehl et al., 2016; Rueckauer et al., 2017; Sengupta et al., 2019) for balancing appropriate ratios between neuronal thresholds and synaptic weights of spiking neuron in the case of ANN-SNN conversion. In this paper, we compare various aspects of our direct-spike trained models with two prior ANN-SNN conversion works (Sengupta et al. 2019; Diehl et al. 2016), which proposed near-lossless ANN-SNN conversion schemes for deep network architectures. The first scheme (Sengupta et al. 2019) balanced the neuronal firing thresholds with respect to corresponding synaptic weights layer-by-layer depending on the actual spiking activities of each layer using a subset of training samples. The second scheme (Diehl et al. 2016) balanced the neuronal firing thresholds with the consideration of ReLU activations in the corresponding ANN layer. Basically, we compare our direct-spike trained model with converted SNNs on the same network architecture in terms of accuracy, inference speed and energy-efficiency. Please note that there are a couple of differences on the network architecture between the conversion networks (Sengupta et al. 2019; Diehl et al. 2016) and our scheme. First, the conversion networks always use average-pooling to reduce the size of previous convolutional output feature-map, whereas our models interchangeably use average pooling or convolve kernels with a stride of 2 in the convolutional layer. Next, the conversion networks only consider identity skip connections for residual SNNs. However, we implement skip connections using either identity mapping or 1×1 convolutional kernel.

3.1.4. Spike Generation Scheme

For the static vision datasets (MNIST, SVHN, and CIFAR-10), each input pixel intensity is converted to a stream of Poisson-distributed spike events that have equivalent firing rates. Specifically, at each time step, the pixel intensity is compared with a uniformly distributed random number (in the range between 0 and 1). If pixel intensity is greater than the random number at the corresponding time step, a spike is generated. This rate-based spike encoding is used to feed the input spikes to the network for a given period of time during both training and inference. For color image datasets, we use the pre-processing technique of horizontal flip before generating input spikes. These input pixels are normalized to represent zero mean and unit standard deviation. Thereafter, we scale the pixel intensities to bound them in the range $[-1,1]$ to represent the whole spectrum of input pixel representations. The normalized pixel intensities are converted to Poisson-distributed spike events such that the generated input signals are bipolar spikes. For the neuromorphic version of the dataset (N-MNIST), we use the original (unfiltered and uncentered) version of spike streams to directly train and test the network in the time domain.

3.1.5. Time-Steps

As mentioned in section 3.1.4, we generate a stochastic Poisson-distributed spike train for each input pixel intensity for event-based operation. The duration of the spike train is very important

TABLE 4 | The deep convolutional spiking neural network architectures for MNIST, N-MNIST, and SVHN dataset.

4 layer network				VGG7				ResNet7			
Layer type	Kernel size	#o/p feature-maps	Stride	Layer type	Kernel size	#o/p feature-maps	Stride	Layer type	Kernel size	#o/p feature-maps	Stride
Convolution	1×5×5	20	1	Convolution	3×3×3	64	1	Convolution	3×3×3	64	1
Average-pooling	2×2		2	Convolution	64×3×3	64	2	Average-pooling	2×2		2
				Average-pooling	2×2		2				
Convolution	20×5×5	50	1	Convolution	64×3×3	128	1	Convolution	64×3×3	128	1
Average-pooling	2×2		2	Convolution	128×3×3	128	2	Convolution	128×3×3	128	2
				Convolution	128×3×3	128	2	Skip convolution	64×1×1	128	2
				Average-pooling	2×2		2				
								Convolution	128×3×3	256	1
								Convolution	256×3×3	256	2
								Skip convolution	128×1×1	256	2
Fully-connected		200		Fully-connected		1024		Fully-connected		1024	
Output		10		Output		10		Output		10	

TABLE 5 | The deep convolutional spiking neural network architectures for a CIFAR-10 dataset.

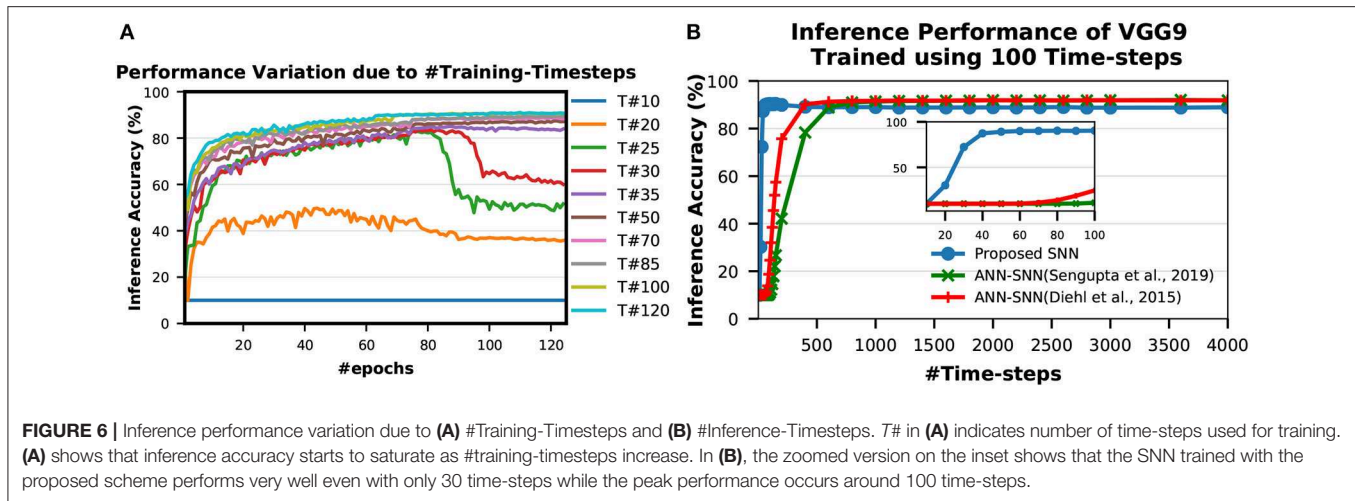
VGG9				ResNet9				ResNet11			
Layer type	Kernel size	#o/p feature-maps	Stride	Layer type	Kernel size	#o/p feature-maps	Stride	Layer type	Kernel size	#o/p feature-maps	Stride
Convolution	3×3×3	64	1	Convolution	3×3×3	64	1	Convolution	3×3×3	64	1
Convolution	64×3×3	64	1	Average-pooling	2×2		2	Average-pooling	2×2		2
Average-pooling	2×2		2								
Convolution	64×3×3	128	1	Convolution	64×3×3	128	1	Convolution	64×3×3	128	1
Convolution	128×3×3	128	1	Convolution	128×3×3	128	1	Convolution	128×3×3	128	1
Average-pooling	2×2		2	Skip convolution	64×1×1	128	1	Skip convolution	64×1×1	128	1
Convolution	128×3×3	256	1	Convolution	128×3×3	256	1	Convolution	128×3×3	256	1
Convolution	256×3×3	256	1	Convolution	256×3×3	256	2	Convolution	256×3×3	256	2
Convolution	256×3×3	256	1	Skip connection	128×1×1	256	2	Skip convolution	128×1×1	256	2
Average-pooling	2×2		2								
				Convolution	256×3×3	512	1	Convolution	256×3×3	512	1
				Convolution	512×3×3	512	2	Convolution	512×3×3	512	1
				Skip convolution	256×1×1	512	2	Skip convolution	512×1×1	512	1
								Convolution	512×3×3	512	1
								Convolution	512×3×3	512	2
								Skip convolution	512×1×1	512	2
Fully-connected		1024		Fully-connected		1024		Fully-connected		1024	
Output		10		Output		10		Output		10	

for SNNs. We measure the length of the spike train (spike time window) in time-steps. For example, a 100 time-step spike train will have approximately 50 random spikes if the corresponding pixel intensity is half in a range of [0,1]. If the number of time-steps (spike time window) is too less, then the SNN will not receive enough information for training or inference. On the other hand, a large number of time-steps will destroy the stochastic property of SNNs and get rid of noise and imprecision at the cost of high latency and power consumption. Hence, the network will not have much energy efficiency over ANN implementations. For these reasons, we experimented with the different number of time-steps to empirically obtain the optimal

number of time-steps required for both training and inference. The experimental process and results are explained in the following subsections.

3.1.5.1. Optimal #time-steps for Training

A spike event can only represent 0 or 1 in each time step, therefore usually its bit precision is considered 1. However, the spike train provides temporal data, which is an additional source of information. Therefore, the spike train length (number of time-steps) in SNN can be considered as its actual precision of neuronal activation. To obtain the optimal #time-steps required for our proposed training method, we trained VGG9 networks on



CIFAR-10 dataset using different time-steps ranging from 10 to 120 (shown in **Figure 6A**). We found that for only 10 time-steps, the network is unable to learn anything as there is not enough information (input precision too low) for the network to be able to learn. This phenomenon is explained by the lack of spikes in the final output. With the initial weights, the accumulated sum of the LIF neuron is not enough to generate output spikes in the later layers. Hence, none of the input spikes propagates to the final output neurons and the output distributions remain 0. Therefore, the computed gradients are always 0 and the network is not updated. For 35–50 time-steps, the network learns well and converges to a reasonable point. From 70 time-steps, the network accuracy starts to saturate. At about 100 time-steps, the network training improvement completely saturates. This is consistent with the bit precision of the inputs. It has been shown in Sarwar et al. (2018) that 8 bit inputs and activations are sufficient to achieve optimal network performance for standard image recognition tasks. Ideally, we need 128 time-steps to represent 8 bit inputs using bipolar spikes. However, 100 time-steps proved to be sufficient as more time-steps provide marginal improvement. We observe a similar trend in VGG7, ResNet7, ResNet9, and ResNet11 SNNs as well while training for SVHN and CIFAR-10 datasets. Therefore, we considered 100 time-steps as the optimal #time-steps for training in our proposed methodology. Moreover, for MNIST dataset, we used 50 time-steps since the required bit precision is only 4 bits (Sarwar et al., 2018).

3.1.5.2. Optimal #time-steps for inference

To obtain the optimal #time-steps required for inferring an image utilizing a network trained with our proposed method, we conducted similar experiments as described in section 3.1.5. We first trained a VGG9 network for CIFAR-10 dataset using 100 time-steps (optimal according to experiments in section 3.1.5). Then, we tested the network performances with different time-steps ranging from 10 to 4,000 (shown in **Figure 6B**). We observed that the network performs very well even with only 30 time-steps while the peak performance occurs around 100

time-steps. For more than 100 time-steps, the accuracy degrades slightly from the peak. This behavior is very different from ANN-SNN converted networks where the accuracy keeps on improving as #time-steps is increased (shown in **Figure 6B**). This can be attributed to the fact that our proposed spike-based training method incorporates the temporal information well into the network training procedure so that the trained network is tailored to perform best at a specific spike time window for the inference. On the other hand, the ANN-SNN conversion schemes are unable to incorporate the temporal information of the input in the trained network. Hence, the ANN-SNN conversion schemes require much higher #time-steps (compared to SNN trained using the proposed method) for the inference in order to resemble input-output mappings similar to ANNs.

3.2. Results

In this section, we analyze the classification performance and efficiency achieved by the proposed spike-based training methodology for deep convolutional SNNs compared to the performance of the transformed SNN using ANN-SNN conversion scheme.

3.2.1. The Classification Performance

Most of the classification performances available in the literature for SNNs are for MNIST and CIFAR-10 datasets. The popular methods for SNN training are “Spike Time Dependent Plasticity (STDP)” based unsupervised learning (Brader et al., 2007; Diehl and Cook, 2015; Srinivasan et al., 2018a,b) and “spike-based backpropagation” based supervised learning (Lee et al., 2016; Neftci et al., 2017; Jin et al., 2018; Mostafa, 2018; Wu et al., 2018b). There are a few works (Kheradpisheh et al., 2016; Tavanaei and Maida, 2016, 2017; Lee et al., 2018) which tried to combine the two approaches to get the best of both worlds. However, these training methods were able to neither train deep SNNs nor achieve good inference performance compared to ANN implementations. Hence, ANN-SNN conversion schemes have been explored by researchers (Cao et al., 2015; Hunsberger and Eliasmith, 2015; Diehl et al., 2016; Rueckauer et al., 2017;

TABLE 6 | Comparison of the SNNs classification accuracies on MNIST, N-MNIST, and CIFAR-10 datasets.

Model	Learning method	Accuracy (MNIST)	Accuracy (N-MNIST)	Accuracy (CIFAR-10)
Hunsberger and Eliasmith (2015)	Offline learning, conversion	98.37%	–	82.95%
Esser et al. (2016)	Offline learning, conversion	–	–	89.32%
Diehl et al. (2016)	Offline learning, conversion	99.10%	–	–
Rueckauer et al. (2017)	Offline learning, conversion	99.44%	–	88.82%
Sengupta et al. (2019)	Offline learning, conversion	–	–	91.55%
Kheradpisheh et al. (2016)	Layerwise STDP + offline SVM classifier	98.40%	–	–
Panda and Roy (2016)	Spike-based autoencoder	99.08%	–	70.16%
Lee et al. (2016)	Spike-based BP	99.31%	98.74%	–
Wu et al. (2018b)	Spike-based BP	99.42%	98.78%	50.70%
Lee et al. (2018)	STDP-based pretraining + spike-based BP	99.28%	–	–
Jin et al. (2018)	Spike-based BP	99.49%	98.88%	–
Wu et al. (2018a)	Spike-based BP	–	99.53%	90.53%
This work	Spike-based BP	99.59%	99.09%	90.95%

Sengupta et al., 2019). Till date, ANN-SNN conversion schemes achieved the best inference performance for CIFAR-10 dataset using deep networks (Rueckauer et al., 2017; Sengupta et al., 2019). Classification performances of all these works are listed in **Table 6** along with ours. To the best of our knowledge, we achieved the best inference accuracy for MNIST using LeNet structured network compared to our spike based training approaches. We also achieved accuracy performance comparable with ANN-SNN converted network (Diehl et al., 2015; Sengupta et al., 2019) for CIFAR-10 dataset while beating all other spike-based training methods.

For a more extensive comparison, we compare the inference performances of trained networks using our proposed methodology with the ANNs and ANN-SNN conversion scheme for same network configuration (depth and structure) side by side in **Table 7**. We also compare with the previous best SNN training results found in the literature that may or may not have the same network depth and structure as ours. The ANN-SNN conversion scheme is a modified and improved version of Sengupta et al. (2019). We are using this modified scheme since it achieves better conversion performance than (Sengupta et al. 2019) as explained in section 3.1.3. Note that all reported classification accuracies are the average of the maximum inference accuracies for 3 independent runs with different seeds.

After initializing the weights, we train the SNNs using a spike-based BP algorithm in an end-to-end manner with Poisson-distributed spike train inputs. Our evaluation of MNIST dataset yields a classification accuracy of 99.59%, which is the best compared to any other SNN training scheme and also identical to other ANN-SNN conversion schemes. We achieve ~96% inference accuracy on SVHN dataset for both trained non-residual and residual SNN. Inference performance for SNNs trained on SVHN dataset has not been reported previously in the literature. We implemented three different networks, as shown in **Table 5**, for classifying CIFAR-10 dataset using a proposed spike-based BP algorithm. For the VGG9 network,

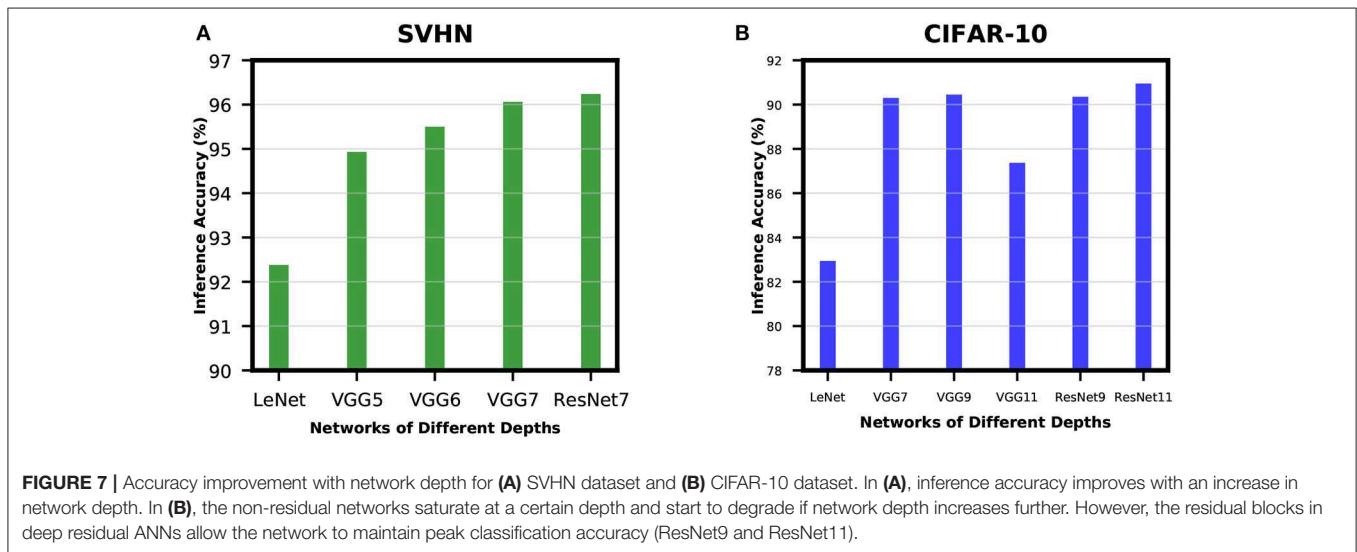
the ANN-SNN conversion schemes provides a near lossless converted network compared to baseline ANN implementation while our proposed training method yields a classification accuracy of 90.45%. For ResNet9 network, the ANN-SNN conversion schemes provide inference accuracy within 0.5–1% of baseline ANN implementation. However, our proposed spike-based training method achieves inference accuracy that is within ~1.5% of baseline ANN implementation. In the case of ResNet11, we observe that the inference accuracy improvement is marginal compared to ResNet9 for baseline ANN implementation and ANN-SNN conversion schemes. However, our proposed SNN training shows improvement of ~0.5% for ResNet11 compared to ResNet9. Overall, our proposed training method achieves comparable inference accuracies for both ResNet and VGG networks compared to baseline ANN implementation and ANN-SNN conversion schemes.

3.2.2. Accuracy Improvement With Network Depth

In order to analyze the effect of network depth for SNNs, we experimented with networks of different depths while training for SVHN and CIFAR-10 datasets. For SVHN dataset, we started with a shallow network derived from LeNet5 model (LeCun et al., 1998) with 2 convolutional and 2 fully-connected layers. This network was able to achieve inference accuracy of only 92.38%. Then, we increased the network depth by adding 1 convolutional layer before the 2 fully-connected layers and we termed this network as VGG5. VGG5 network was able to achieve significant improvement over its predecessor. Similarly, we tried VGG6 followed by VGG7, and the improvement started to become very small. We have also trained ResNet7 to understand how residual networks perform compared to non-residual networks of similar depth. The results of these experiments are shown in **Figure 7A**. We carried out similar experiments for CIFAR-10 dataset as well. The results show a similar trend as described in **Figure 7B**. These results ensure that network depth improves the learning capacity of direct-spike trained SNNs similar to ANNs. The non-residual networks saturate at a certain depth and start to degrade

TABLE 7 | Comparison of classification performance.

Dataset	Model	Inference Accuracy				
		ANN	ANN-SNN (Diehl et al., 2015)	ANN-SNN (Sengupta et al., 2019)	SNN (Previous Best)	SNN This Work
MNIST	LeNet	99.57%	99.55%	99.59%	99.49% (Jin et al., 2018)	99.59%
N-MNIST	LeNet	–	–	–	99.53% (Wu et al., 2018a)	99.09%
SVHN	VGG7	96.36%	96.33%	96.30%	–	96.06%
	ResNet7	96.43%	96.33%	96.40%	–	96.21%
CIFAR-10	VGG9	91.98%	91.89%	92.01%	–	90.45%
	ResNet9	91.85%	90.78%	91.59%	90.53% (Wu et al., 2018a)	90.35%
	ResNet11	91.87%	90.98%	91.65%	–	90.95%



if network depth is further increased (VGG11 in **Figure 7B**) due to the degradation problem mentioned in He et al. (2016). In such a scenario, the residual connections in deep residual ANNs allow the network to maintain peak classification accuracy utilizing the skip connections (He et al., 2016), as seen in **Figure 7B** (ResNet9 and ResNet11).

4. DISCUSSION

4.1. Comparison With Relevant Works

In this section, we compare our proposed supervised learning algorithm with other recent spike-based BP algorithms. The spike-based learning rules primarily focus on directly training and testing SNNs with spike-trains, and no conversion is necessary for applying in real-world spiking scenario. In recent years, there are an increasing number of supervised gradient descent method in spike-based learning. The Panda and Roy (2016) developed a spike-based auto-encoder mechanism to train deep convolutional SNNs. They dealt with membrane potential as a differentiable signal and showed recognition capabilities in standard vision tasks (MNIST and CIFAR-10 datasets). Meanwhile, Lee et al. (2016) followed the approach

using differentiable membrane potential to explore a spike-based BP algorithm in an end-to-end manner. In addition, Lee et al. (2016) presented the error normalization scheme to prevent exploding gradient phenomena while training deep SNNs. Researchers in Jin et al. (2018) proposed hybrid macro/micro level backpropagation (HM2-BP). HM2-BP is developed to capture the temporal effect of the individual spike (in micro-level) and rate-encoded error (at macro-level). The reference Shrestha and Orchard (2018) employed exponential function for the approximate derivative of neuronal function and developed a credit assignment scheme to calculate the temporal dependencies of error throughout the layers. Huh and Sejnowski (2018) has trained recurrent spiking networks by replacing the threshold with a gate function and employing BPTT technique (Werbos, 1990). While BPTT technique has been a popular method to train recurrent artificial and spiking recurrent networks, Lillicrap and Santoro (2019) points out the storing and retrieving past variables and differentiation them through time in biological neurons seems to be impossible. Recently, e-prop Bellec et al. (2019) presented an approximation method to bypass neuronal state savings for enhancing the computational efficiency of BPTT. In temporal spike encoding domain, Mostafa (2018) proposed an interesting temporal

spike-based BP algorithm by treating the spike-time as the differential activation of neuron. Temporal encoding based SNN has the potential to process the tasks with the small number of spikes. All of these works demonstrated spike-based learning in simple network architectures and has a large gap in classification accuracy compared to deep ANNs. More recently, Wu et al. (2018a) presented a neuron normalization technique (called NeuNorm) that calculates the average input firing rates to adjust neuron selectivity. NeuNorm enables spike-based training within a relatively short time-window while achieving competitive performances. In addition, they presented an input encoding scheme that receives both spike and non-spike signals for preserving the precision of input data.

There are several points that distinguish our work from others. First, we use a pseudo derivative method that accounts for leaky effect in membrane potential of LIF neurons. We approximately estimate the leaky effect by comparing total membrane potential value and obtain the ratio between IF and LIF neurons. During the back-propagating phase, the pseudo derivative of LIF neuronal function is estimated by combining the straight through estimation and leak correctional term as described in Equation (22). Next, we construct our networks by leveraging frequently used architectures such as VGG (Simonyan and Zisserman, 2014) and ResNet (He et al., 2016). To the best of our knowledge, this is the first work that demonstrates spike-based supervised BP learning for SNNs containing more than 10 trainable layers. Our deep SNNs obtain the superior classification accuracies in MNIST, SVHN, and CIFAR-10 datasets in comparison to the other networks trained with the spike-based algorithm. In addition, as opposed to complex error or neuron normalization method adopted by Lee et al. (2016) and Wu et al. (2018a), respectively, we demonstrate that deep SNNs can be naturally trained by only accounting for spiking activities of the network. As a result, our work paves the effective way for training deep SNNs with a spike-based BP algorithm.

4.2. Spike Activity Analysis

The most important advantage of event-based operation of neural networks is that the events are very sparse in nature. To verify this claim, we analyzed the spiking activities of the direct-spike trained SNNs and ANN-SNN converted networks in the following subsections.

4.2.1. Spike Activity per Layer

The layer-wise spike activities of both SNN trained using our proposed methodology, and ANN-SNN converted network (using scheme 1) for VGG9 and ResNet9 are shown in **Figures 8A,B**, respectively. In the case of ResNet9, only the first average pooling layer's output spike activity is shown in the figure as for the direct-spike trained SNN, the other spatial-poolings are done by stride 2 convolutions. In **Figure 8**, it can be seen that the input layer has the highest spike activity that is significantly higher than any other layer. The spike activity reduces significantly as the network depth increases.

We can observe from **Figures 8A,B** that the average spike activity in a direct-spike trained SNN is much higher than ANN-SNN converted network. The ANN-SNN converted network uses

a higher threshold compared to 1 (in case of direct-spike trained SNN) since the conversion scheme applies layer-wise neuronal threshold modulation. This higher threshold reduces spike activity in ANN-SNN converted networks. However, in both cases, the spike activity decreases with increasing network depth.

4.2.2. #Spikes/Inference

From **Figure 8**, it is evident that average spike activity in ANN-SNN converted networks is much less than in the direct-spike trained SNN. However, for inference, the network has to be evaluated over many time-steps. Therefore, to quantify the actual spike activity for an inference operation, we measured the average number of spikes required for inferring one image. For this purpose, we counted the number of spikes generated (including input spikes) for classifying the test set of a particular dataset for a specific number of time-steps and averaged the count for generating the quantity “#spikes per image inference.” We have used two different time-steps for ANN-SNN converted networks; one for iso-accuracy comparison and the other for maximum accuracy comparison with the direct-spike trained SNNs. Iso-accuracy inference requires less #time-steps than maximum accuracy inference, hence has a lower number of spikes per image inference. For few networks, the ANN-SNN conversion scheme always provides accuracy less than or equal to the direct-spike trained SNN. Hence, we only compare spikes per image inference in maximum accuracy condition for those ANN-SNN converted networks while comparing with direct-spike trained SNNs. For the analysis, we quantify the spike-efficiency (amount reduction in #spikes) from the #spikes/image inference. The results are listed in **Table 8**, where the 1st row corresponds to iso-accuracy and the 2nd row corresponds to maximum-accuracy condition for each network. As shown in **Table 8**, the direct-spike trained SNNs are more efficient in terms of #spikes/inference compared to the ANN-SNN converted networks for the maximum accuracy condition. For an iso-accuracy condition, only deep SNNs (such as VGG9 and ResNet11) are more efficient in terms of #spikes/inference compared to the ANN-SNN converted networks.

The **Figure 9** shows the relationship between inference accuracy, latency and #spikes/inference for ResNet11 networks trained on CIFAR-10 dataset. We can observe that #spikes/inference is higher for direct-spike trained SNN compared to ANN-SNN converted networks at any particular latency. However, SNN trained with spike-based BP requires only 100 time-steps for maximum inference accuracy, whereas ANN-SNN converted networks require 3,000–3,500 time-steps to reach maximum inference accuracy. Hence, under maximum accuracy condition, direct-spike trained ResNet11 requires much fewer #spikes/inference compared to ANN-SNN converted networks, while achieving comparable accuracy. Even under iso-accuracy condition, the direct-spike trained ResNet11 requires fewer #spikes/inference compared to the ANN-SNN converted networks (**Table 8**).

4.3. Inference Speedup

The time required for inference is linearly proportional to the #time-steps (**Figure 9**). Hence, we can also quantify the inference speedup for direct-spike trained SNNs compared to ANN-SNN

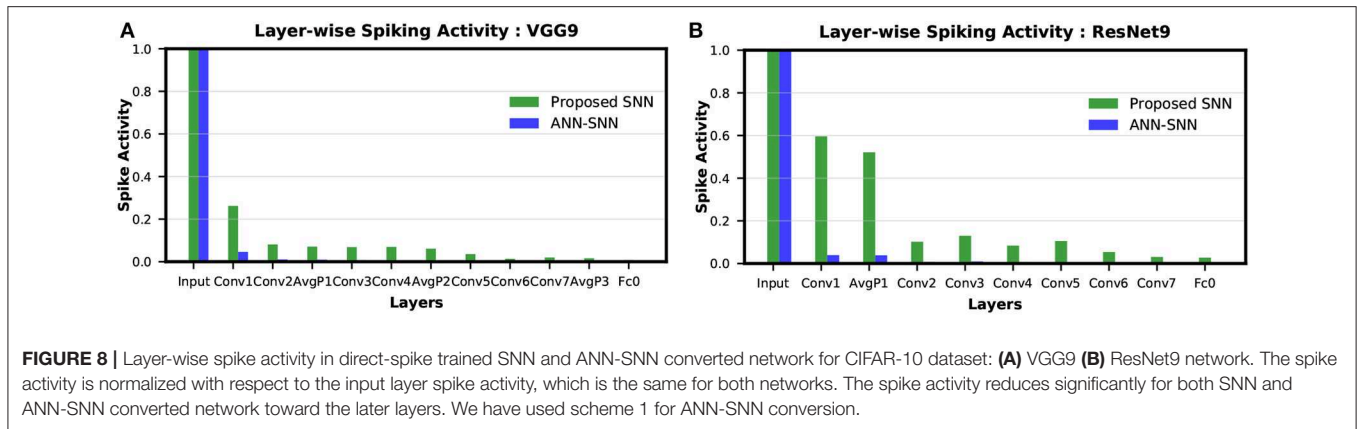


TABLE 8 | #Spikes/Image inference and spike efficiency comparison between SNN and ANN-SNN converted networks for benchmark datasets trained on different network models.

Dataset	Model	Spike/Image			Spike efficiency compared to	
		SNN	ANN-SNN (Diehl et al., 2015)	ANN-SNN (Sengupta et al., 2019)	ANN-SNN (Diehl et al., 2015)	ANN-SNN (Sengupta et al., 2019)
MNIST	LeNet	5.52E+04	3.4E+04	2.9E+04	0.62x	0.53x
				7.3E+04		1.32x
SVHN	VGG7	5.56E+06	3.7E+06	1.0E+07	0.67x	1.84x
	ResNet7		1.9E+07	1.7E+07	3.40x	2.99x
			4.66E+06	3.9E+06	3.1E+06	0.85x
	VGG9	1.24E+06	1.6E+06	2.2E+06	1.32x	1.80x
CIFAR-10	ResNet9		8.3E+06	9.6E+06	6.68x	7.78x
			4.32E+06	2.7E+06	1.5E+06	0.63x
	ResNet11		1.0E+07	7.8E+06	2.39x	1.80x
			1.53E+06	9.7E+06	1.8E+06	6.33x
			9.2E+06		5.99x	

(For each network, the 1st row corresponds to iso-accuracy and the 2nd row corresponds to maximum-accuracy condition).

converted networks from the #time-steps required for inference, as shown in **Table 9**. For example, for VGG9 network, the proposed training method can achieve 8x (5x) speedup for iso-accuracy and up to 36x (25x) speedup for maximum accuracy in inference compared to respective ANN-SNN converted networks [i.e., scheme 1 (Sengupta et al., 2019) and scheme 2 (Diehl et al., 2016)]. Similarly, for ResNet networks, the proposed training method can achieve 6x speedup for iso-accuracy and up to 35x speedup for maximum accuracy condition in inference. It is interesting to note that direct-spike trained SNN is always more efficient in terms of time-steps compared to the equivalent ANN-SNN conversion network, but not in terms of the number of spikes, in some cases. It will require a detailed investigation to determine if ANN-SNN methods used higher firing rates, whether they would be able to classify quickly as well, while incurring a lower number of spike/inference.

4.4. Complexity Reduction

Deep ANNs struggle to meet the demand of extraordinary computational requirements. SNNs can mitigate this effort by

enabling efficient event-based computations. To compare the computational complexity of these two cases, we first need to understand the operation principle of both. An ANN operation for inferring the category of a particular input requires a single feed-forward pass per image. For the same task, the network must be evaluated over a number of time-steps in the spiking domain. If regular hardware is used for both ANN and SNN, then it is evident that SNN will have computation complexity in the order of hundreds or thousands more compared to an ANN. However, there are specialized hardwares that account for the event-based neural operation and “computes only when required” for inference. SNNs can potentially exploit such alternative mechanisms of network operation and carry out an inference operation in the spiking domain much more efficiently than an ANN. Also, for deep SNNs, we have observed the increase in sparsity as the network depth increases. Hence, the benefits from event-based neuromorphic hardware are expected to increase as the network depth increases.

An estimate of the actual energy consumption of SNNs and comparison with ANNs is outside the scope of this

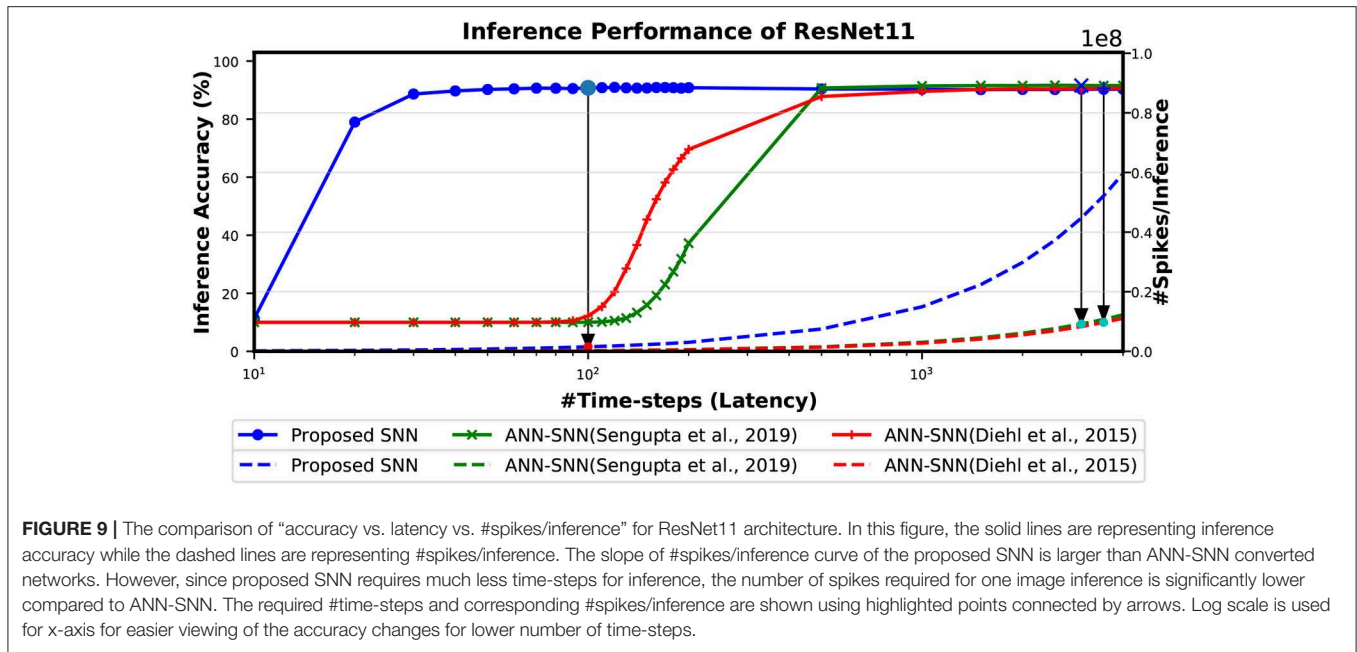


FIGURE 9 | The comparison of “accuracy vs. latency vs. #spikes/inference” for ResNet11 architecture. In this figure, the solid lines are representing inference accuracy while the dashed lines are representing #spikes/inference. The slope of #spikes/inference curve of the proposed SNN is larger than ANN-SNN converted networks. However, since proposed SNN requires much less time-steps for inference, the number of spikes required for one image inference is significantly lower compared to ANN-SNN. The required #time-steps and corresponding #spikes/inference are shown using highlighted points connected by arrows. Log scale is used for x-axis for easier viewing of the accuracy changes for lower number of time-steps.

TABLE 9 | Inference #time-steps and corresponding speedup comparison between SNN and ANN-SNN converted networks for benchmark datasets trained on different network models.

Dataset	Model	Timesteps			SNN Inference Speedup Compared to	
		SNN	ANN-SNN (Diehl et al., 2015)	ANN-SNN (Sengupta et al., 2019)	ANN-SNN (Diehl et al., 2015)	ANN-SNN (Sengupta et al., 2019)
MNIST	LeNet	50	180	200	3.6x	4x
				500		10x
SVHN	VGG7	100	500	1,600	5x	16x
			2,500	2,600	25x	26x
	ResNet7	100	500	400	5x	4x
			3,000	2,500	30x	25x
CIFAR-10	VGG9	100	500	800	5x	8x
			2,500	3,600	25x	36x
	ResNet9	100	800	600	8x	6x
			3,000	3,000	30x	30x
ResNet11	100	3500	600	35x	6x	
			3,000	30x		

(For each network, the 1st row corresponds to iso-accuracy and the 2nd row corresponds to maximum-accuracy condition).

work. However, we can gain some insight by quantifying the computational energy consumption for a synaptic operation and comparing the number of synaptic operations being performed in the ANN vs. the SNN trained with our proposed algorithm and ANN-SNN converted network. We can estimate the number of synaptic operations per layer of a neural network from the structure for the convolutional and linear layers. In an ANN, a multiply-accumulate (MAC) computation is performed per synaptic operation. While a specialized SNN hardware would perform simply an accumulate computation (AC) per synaptic operation only if an incoming spike is received. Hence, the total number of AC operations in a SNN can be estimated by the

layer-wise product and summation of the average neural spike count for a particular layer and the corresponding number of synaptic connections. We also have to multiply the #time-steps with the #AC operations to get total #AC operation for one inference. For example, assume that there are L layers each with N_l neurons, S_l synaptic connections and a_l average spiking activity where l is the layer number. Then, the total number of synaptic operations in a layer is $N_l \times S_l \times a_l$. The $N_l \times S_l$ is equal to the ANN (#MAC) operations of a particular layer. Therefore, the total number of synaptic operations in a layer of an SNN becomes $\#MAC_l \times a_l$. The total number of #AC operations required for an image inference is the sum of synaptic

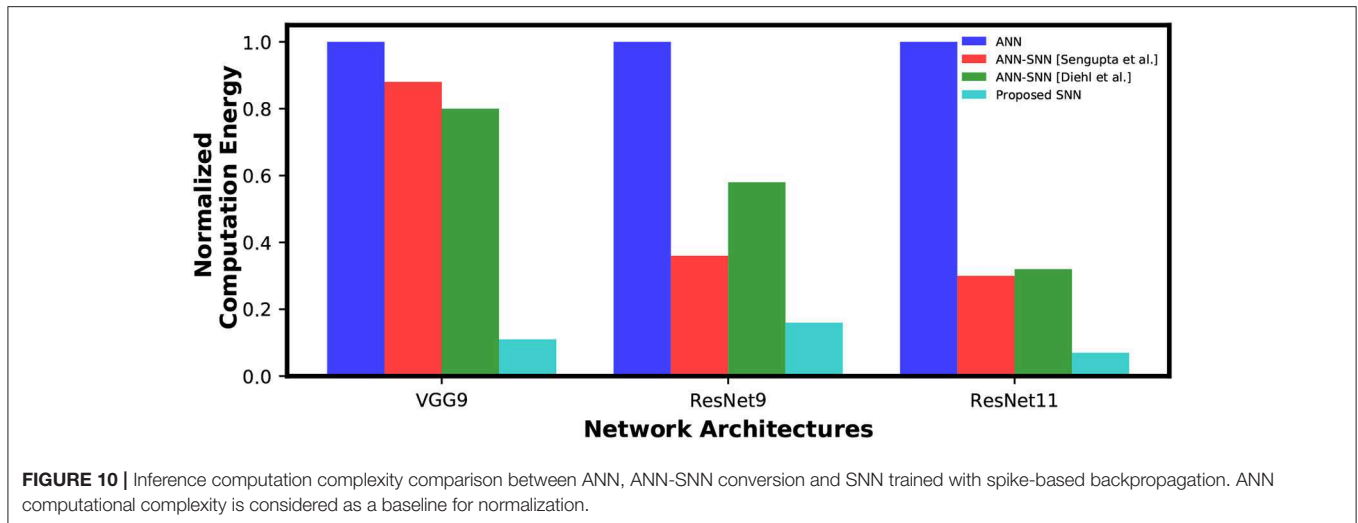


FIGURE 10 | Inference computation complexity comparison between ANN, ANN-SNN conversion and SNN trained with spike-based backpropagation. ANN computational complexity is considered as a baseline for normalization.

TABLE 10 | Iso-spike comparison for optimal condition.

Dataset	Model	Time-steps			Accuracy (%)		
		SNN	ANN-SNN		SNN	ANN-SNN	
			(Diehl et al., 2015)	(Sengupta et al., 2019)		(Diehl et al., 2015)	(Sengupta et al., 2019)
MNIST	LeNet	20	62	75	99.36	99.19	88.62
SVHN	VGG7	30	235	235	95.00	95.34	88.13
	ResNet7	30	200	200	95.06	95.63	95.48
	VGG9	50	228	260	89.33	69.53	61.08
CIFAR-10	ResNet9	50	390	490	89.52	89.51	90.06
	ResNet11	50	307	280	90.24	82.75	73.82

SNN time-steps corresponds to the latency to reach accuracy within ~1% of maximum accuracy. ANN-SNN time-steps corresponds to the latency required for same number of spike/inference as SNN to occur. SNN and ANN-SNN accuracies are accuracies corresponding to respective latency.

operations in all layers during the inference time-window. Hence, $\#AC/inference = (\sum_l (\#MAC_l \times a_l)) \times \#timesteps$. This formula is used for estimating both ANN-SNN AC operations and SNN AC operations per image inference. On the other hand, the number of ANN (MAC) operation per inference becomes simply, $\#MAC/inference = \sum_l (\#MAC_l)$. Based on this concept, we estimated the total number of MAC operations for ANN, and the total number of AC operations for direct-spike trained SNN and ANN-SNN converted network, for VGG9, ResNet9 and ResNet11. The ratio of ANN-SNN converted networks' (scheme1-scheme2) AC operations to direct-spike trained SNN AC operations to ANN MAC operations is (28.18–25.60):3.61:1 for VGG9 while the ratio is (11.67–18.42):5.06:1 for the ResNet9 and (9.6–10.16):2.09:1 for ResNet11 (for maximum accuracy condition).

However, a MAC operation usually consumes an order of magnitude more energy than an AC operation. For instance, according to Han et al. (2015), a 32-bit floating point MAC operation consumes 4.6 pJ and a 32-bit floating point AC operation consumes 0.9 pJ in 45 nm technology node. Hence, one synaptic operation in an ANN is equivalent to ~5 synaptic operations in a SNN. Moreover, 32-bit floating point

computation can be replaced by fixed point computation using integer MAC and AC units without losing accuracy since the conversion is reported to be almost loss-less Lin et al. (2016). A 32-bit integer MAC consumes roughly 3.2 pJ while a 32-bit AC operation consumes only 0.1 pJ in 45nm process technology. Considering this fact, our calculations demonstrate that the SNNs trained using the proposed method will be 7.81x~7.10x and 8.87x more computationally energy-efficient for inference compared to the ANN-SNN converted networks and an ANN, respectively, for the VGG9 network architecture. We also gain 4.6x~4.87x(2.31x~3.64x) and 15.32x(6.32x) energy-efficiency, for the ResNet11(ResNet9) network, compared to the ANN-SNN converted networks and an ANN, respectively. The **Figure 10** shows the reduction in computation complexity for ANN-SNN conversions and SNN trained with the proposed methodology compared to ANNs.

It is worth noting here that as the sparsity of the spike signals increases with an increase in network depth in SNNs. Hence, the energy-efficiency is expected to increase almost exponentially in both ANN-SNN conversion network (Sengupta et al., 2019) and SNN trained with proposed methodology compared to an ANN implementation. The depth of network is the key factor

for achieving a significant increase in the energy efficiency for SNNs in neuromorphic hardware. However, the computational efficiency does not perfectly align with the overall efficiency since the dominant energy consumption can be the memory traffic on von-Neumann computing hardware. The dataflows in asynchronous SNNs are less predictable and more complicated. Hence, a detailed study is required to estimate the overall efficiency of SNNs accurately.

4.5. Iso-Spike Comparison for Optimal Condition

In section 4.3, we observe that SNNs trained with proposed method achieve significant speed-up in both max-accuracy and iso-accuracy condition. However, in section 4.2.2, we found that the proposed method is in some cases (in an iso-accuracy condition) not more efficient than ANN-SNN conversions in terms of #spike/inference. The reason behind it is that an iso-accuracy condition may not be optimal for the SNNs trained with proposed method. In an iso-accuracy case, we have used max-accuracy latency (50 time-steps for MNIST and 100 time-steps for other networks) for direct-spike trained SNN, whereas most of the conversion networks used much less latency than the max-accuracy condition. In view of this, there is a need to determine the circumstances where our proposed method performs as well as or better than the SNN-ANN conversion methods on spike count, time steps, and accuracy. Consequently, in this section we analyze another interesting comparison.

In this analysis, we compare our proposed method and ANN-SNN conversion methods (Diehl et al., 2015; Sengupta et al., 2019) under the optimal condition at equal number of spikes. We define the “optimal #time-steps” for SNNs trained with our proposed method as the #time-steps required to reach within ~1% of peak accuracy (when the accuracy starts to saturate). Based on this definition, we observed that the optimal #time-steps for MNIST, SVHN, CIFAR10 networks are 20, 30, and 50, respectively. For this comparison, we recorded the achieved accuracy and #spike/inference of the SNNs trained with our proposed method for the corresponding optimal #time-steps. Then, we ran ANN-SNN networks for a length of time such that they use the similar number of spikes. In this iso-spike condition, we recorded the accuracy of the ANN-SNN networks (for both conversion methods) and the number of time-steps they require. The results are summarized in **Table 10**.

For comparatively shallower networks such as LeNet, VGG7 (VGG type) and ResNet7, ResNet9 (Residual type), the ANN-SNN conversion networks achieve as good as or slightly better accuracy at iso-spike condition compared to the SNNs trained with our proposed method. However, these ANN-SNN conversion networks require 3x-10x higher latency for inference. On the other hand, for deeper networks such as VGG9 and ResNet11, the ANN-SNN conversion networks achieve significantly lower accuracy compared to SNNs trained with our proposed method even with much higher latency. This trend indicates that for deeper networks, SNNs trained with

our proposed method will be more energy-efficient than the conversion networks under an iso-spike condition.

5. CONCLUSION

In this work, we propose a spike-based backpropagation training methodology for popular deep SNN architectures. This methodology enables deep SNNs to achieve comparable classification accuracies on standard image recognition tasks. Our experiments show the effectiveness of the proposed learning strategy on deeper SNNs (7–11 layer VGG and ResNet network architectures) by achieving the best classification accuracies in MNIST, SVHN, and CIFAR-10 datasets among other networks trained with spike-based learning till date. The performance gap in terms of quality between ANN and SNN is substantially reduced by the application of our proposed methodology. Moreover, significant computational energy savings are expected when deep SNNs (trained with the proposed method) are employed on suitable neuromorphic hardware for the inference.

DATA AVAILABILITY STATEMENT

Publicly available datasets were analyzed in this study. This data can be found here: MNIST, N-MNIST, SVHN, CIFAR-10 datasets. The source code is publicly released at “https://github.com/chan8972/Enabling_Spikebased_Backpropagation”.

AUTHOR CONTRIBUTIONS

CL and SS implemented the algorithm and conducted the experiments. CL, SS, PP, GS, and KR discussed about the results and analysis, and wrote the manuscript.

FUNDING

This work was supported in part by C-BRIC, one of six centers in JUMP, a Semiconductor Research Corporation (SRC) program sponsored by DARPA, the National Science Foundation, Intel Corporation, the DoD Vannevar Bush Fellowship and the U.S. Army Research Laboratory and the U.K. Ministry of Defence under Agreement Number W911NF-16-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

ACKNOWLEDGMENTS

This manuscript has been released as a Pre-Print at Lee et al. (2019a). We would like to thank Dr. Gerard (Rod) Rinkus for helpful comments.

REFERENCES

- Ankit, A., Sengupta, A., Panda, P., and Roy, K. (2017). "Resparc: a reconfigurable and energy-efficient architecture with memristive crossbars for deep spiking neural networks," in *Proceedings of the 54th Annual Design Automation Conference 2017* (New York, NY), 1–6.
- Bellec, G., Scherr, F., Subramoney, A., Hajek, E., Salaj, D., Legenstein, R., et al. (2019). A solution to the learning dilemma for recurrent networks of spiking neurons. *bioRxiv [Preprint]*. doi: 10.1101/738385
- Bengio, Y., Léonard, N., and Courville, A. (2013). Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv [Preprint]*. arXiv:1308.3432.
- Bohte, S. M., Kok, J. N., and La Poutre, H. (2002). Error-backpropagation in temporally encoded networks of spiking neurons. *Neurocomputing* 48, 17–37. doi: 10.1016/S0925-2312(01)00658-0
- Brader, J. M., Senn, W., and Fusi, S. (2007). Learning real-world stimuli in a neural network with spike-driven synaptic dynamics. *Neural Comput.* 19, 2881–2912. doi: 10.1162/neco.2007.19.11.2881
- Brette, R., and Gerstner, W. (2005). Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. *J. Neurophysiol.* 94, 3637–3642. doi: 10.1152/jn.00686.2005
- Cao, Y., Chen, Y., and Khosla, D. (2015). Spiking deep convolutional neural networks for energy-efficient object recognition. *Int. J. Comput. Vision* 113, 54–66. doi: 10.1007/s11263-014-0788-3
- Davies, M., Srinivasa, N., Lin, T.-H., China, G., Cao, Y., Choday, S. H., et al. (2018). Loihi: a neuromorphic manycore processor with on-chip learning. *IEEE Micro* 38, 82–99. doi: 10.1109/MM.2018.112130359
- Dayan, P., and Abbott, L. F. (2001). *Theoretical Neuroscience, Vol. 806*. Cambridge, MA: MIT Press.
- Diehl, P. U., and Cook, M. (2015). Unsupervised learning of digit recognition using spike-timing-dependent plasticity. *Front. Comput. Neurosci.* 9:99. doi: 10.3389/fncom.2015.00099
- Diehl, P. U., Neil, D., Binas, J., Cook, M., Liu, S.-C., and Pfeiffer, M. (2015). "Fast-classifying, high-accuracy spiking deep networks through weight and threshold balancing," in *2015 International Joint Conference on Neural Networks (IJCNN)* (Killarney: IEEE), 1–8.
- Diehl, P. U., Zarella, G., Cassidy, A., Pedroni, B. U., and Neftci, E. (2016). "Conversion of artificial recurrent neural networks to spiking neural networks for low-power neuromorphic hardware," in *2016 IEEE International Conference on Rebooting Computing (ICRC)* (San Diego, CA: IEEE), 1–8.
- Esser, S., Merolla, P., Arthur, J., Cassidy, A., Appuswamy, R., Andreopoulos, A., et al. (2016). Convolutional networks for fast, energy-efficient neuromorphic computing. *arXiv [Preprint]*. Available online at: <http://arxiv.org/abs/1603.08270>
- Furber, S. B., Lester, D. R., Plana, L. A., Garside, J. D., Painkras, E., Temple, S., et al. (2013). Overview of the spinnaker system architecture. *IEEE Trans. Comput.* 62, 2454–2467. doi: 10.1109/TC.2012.142
- Han, S., Pool, J., Tran, J., and Dally, W. (2015). "Learning both weights and connections for efficient neural network," in *Advances in Neural Information Processing Systems* (Montréal, QC), 1135–1143.
- He, K., Zhang, X., Ren, S., and Sun, J. (2015). "Delving deep into rectifiers: surpassing human-level performance on imagenet classification," in *Proceedings of the IEEE International Conference on Computer Vision* (Santiago), 1026–1034.
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). "Deep residual learning for image recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Las Vegas, NV), 770–778.
- Hodgkin, A. L., and Huxley, A. F. (1952). Currents carried by sodium and potassium ions through the membrane of the giant axon of lolo. *J. Physiol.* 116, 449–472. doi: 10.1113/jphysiol.1952.sp004717
- Huh, D., and Sejnowski, T. J. (2018). "Gradient descent for spiking neural networks," in *Advances in Neural Information Processing Systems* (Montréal, QC), 1440–1450.
- Hunsberger, E., and Eliasmith, C. (2015). Spiking deep networks with lif neurons. *arXiv [Preprint]*. arXiv:1510.08829.
- Ioffe, S., and Szegedy, C. (2015). Batch normalization: accelerating deep network training by reducing internal covariate shift. *arXiv [Preprint]*. arXiv:1502.03167.
- Izhikevich, E. M. (2003). Simple model of spiking neurons. *IEEE Trans. Neural Netw.* 14, 1569–1572. doi: 10.1109/TNN.2003.820440
- Jin, Y., Li, P., and Zhang, W. (2018). Hybrid macro/micro level backpropagation for training deep spiking neural networks. *arXiv [Preprint]*. arXiv:1805.07866.
- Kappel, D., Habenschuss, S., Legenstein, R., and Maass, W. (2015). Network plasticity as bayesian inference. *PLoS Comput. Biol.* 11:e1004485. doi: 10.1371/journal.pcbi.1004485
- Kappel, D., Legenstein, R., Habenschuss, S., Hsieh, M., and Maass, W. (2018). A dynamic connectome supports the emergence of stable computational function of neural circuits through reward-based learning. *Eneuro* 5:ENEURO.0301-17.2018. doi: 10.1523/ENEURO.0301-17.2018
- Kheradpisheh, S. R., Ganjtabesh, M., Thorpe, S. J., and Masquelier, T. (2016). Stdp-based spiking deep neural networks for object recognition. *arXiv [Preprint]*. arXiv:1611.01421.
- Kingma, D. P., and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv [Preprint]*. arXiv:1412.6980.
- Krizhevsky, A., and Hinton, G. (2009). *Learning Multiple Layers of Features From Tiny Images*. Technical report, Citeseer.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. *Adv. Neural. Inform. Process. Syst.* 25, 1097–1105.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 2278–2324. doi: 10.1109/5.726791
- Lee, C., Panda, P., Srinivasan, G., and Roy, K. (2018). Training deep spiking convolutional neural networks with stdp-based unsupervised pre-training followed by supervised fine-tuning. *Front. Neurosci.* 12:435. doi: 10.3389/fnins.2018.00435
- Lee, C., Sarwar, S. S., and Roy, K. (2019a). Enabling spike-based backpropagation in state-of-the-art deep neural network architectures. *arXiv [Preprint]*. arXiv:1903.06379.
- Lee, C., Srinivasan, G., Panda, P., and Roy, K. (2019b). Deep spiking convolutional neural network trained with unsupervised spike-timing-dependent plasticity. *IEEE Trans. Cogn. Dev. Syst.* 11, 384–394. doi: 10.1109/TCDS.2018.2833071
- Lee, J. H., Delbruck, T., and Pfeiffer, M. (2016). Training deep spiking neural networks using backpropagation. *Front. Neurosci.* 10:508. doi: 10.3389/fnins.2016.00508
- Lichtsteiner, P., Posch, C., and Delbruck, T. (2008). A 128 × 128 120 db 15 μs latency asynchronous temporal contrast vision sensor. *IEEE J. Solid State Circuits* 43, 566–576. doi: 10.1109/JSSC.2007.914337
- Lillicrap, T. P., and Santoro, A. (2019). Backpropagation through time and the brain. *Curr. Opin. Neurobiol.* 55, 82–89. doi: 10.1016/j.conb.2019.01.011
- Lin, D., Talathi, S., and Annapureddy, S. (2016). "Fixed point quantization of deep convolutional networks," in *International Conference on Machine Learning* (New York, NY), 2849–2858.
- Maass, W. (1997). Networks of spiking neurons: the third generation of neural network models. *Neural Netw.* 10, 1659–1671. doi: 10.1016/S0893-6080(97)00011-7
- Merolla, P. A., Arthur, J. V., Alvarez-Icaza, R., Cassidy, A. S., Sawada, J., Akopyan, F., et al. (2014). A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science* 345, 668–673. doi: 10.1126/science.1254642
- Mostafa, H. (2018). Supervised learning based on temporal coding in spiking neural networks. *IEEE Trans. Neural Netw. Learn. Syst.* 29, 3227–3235. doi: 10.1109/TNNLS.2017.2726060
- Neftci, E. O., Augustine, C., Paul, S., and Detorakis, G. (2017). Event-driven random back-propagation: enabling neuromorphic deep learning machines. *Front. Neurosci.* 11:324. doi: 10.3389/fnins.2017.00324
- Neftci, E. O., Pedroni, B. U., Joshi, S., Al-Shedivat, M., and Cauwenberghs, G. (2015). Unsupervised learning in synaptic sampling machines. *arXiv [Preprint]*. arXiv:1511.04484.
- Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., and Ng, A. Y. (2011). "Reading digits in natural images with unsupervised feature learning," in *NIPS Workshop on Deep Learning and Unsupervised Feature Learning, Vol. 2011*, 5.
- Orchard, G., Jayawant, A., Cohen, G. K., and Thakor, N. (2015). Converting static image datasets to spiking neuromorphic datasets using saccades. *Front. Neurosci.* 9:437. doi: 10.3389/fnins.2015.00437

- Panda, P., and Roy, K. (2016). "Unsupervised regenerative learning of hierarchical features in spiking deep networks for object recognition," in *2016 International Joint Conference on Neural Networks (IJCNN)* (Vancouver, BC: IEEE), 299–306.
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., et al. (2017). "Automatic differentiation in pytorch," in *NIPS 2017 Autodiff Workshop: The Future of Gradient-based Machine Learning Software and Techniques* (Long Beach, CA).
- Rueckauer, B., Lungu, I.-A., Hu, Y., Pfeiffer, M., and Liu, S.-C. (2017). Conversion of continuous-valued deep networks to efficient event-driven networks for image classification. *Front. Neurosci.* 11:682. doi: 10.3389/fnins.2017.00682
- Rumelhart, D. E., Hinton, G. E., and Williams, R. J. (1985). *Learning Internal Representations by Error Propagation*. Technical report, California Univ San Diego La Jolla Inst for Cognitive Science.
- Sarwar, S. S., Srinivasan, G., Han, B., Wijesinghe, P., Jaiswal, A., Panda, P., et al. (2018). Energy efficient neural computing: a study of cross-layer approximations. *IEEE J. Emerg. Sel. Top. Circuits Syst.* 8, 796–809. doi: 10.1109/JETCAS.2018.2835809
- Sengupta, A., Ye, Y., Wang, R., Liu, C., and Roy, K. (2019). Going deeper in spiking neural networks: Vgg and residual architectures. *Front. Neurosci.* 13:95. doi: 10.3389/fnins.2019.00095
- Shrestha, S. B., and Orchard, G. (2018). "Slayer: spike layer error reassignment in time," in *Advances in Neural Information Processing Systems* (Montréal, QC), 1412–1421.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., et al. (2016). Mastering the game of go with deep neural networks and tree search. *Nature* 529:484. doi: 10.1038/nature16961
- Simonyan, K., and Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. *arXiv [Preprint]*. arXiv:1409.1556.
- Srinivasan, G., Panda, P., and Roy, K. (2018a). Spilinc: Spiking liquid-ensemble computing for unsupervised speech and image recognition. *Front. Neurosci.* 12:524. doi: 10.3389/fnins.2018.00524
- Srinivasan, G., Panda, P., and Roy, K. (2018b). Stdp-based unsupervised feature learning using convolution-over-time in spiking neural networks for energy-efficient neuromorphic computing. *ACM J. Emerg. Technol. Comput. Syst.* 14:44. doi: 10.1145/3266229
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. (2014). Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* 15, 1929–1958. doi: 10.5555/2627435.2670313
- Tavanaei, A., and Maida, A. S. (2016). Bio-inspired spiking convolutional neural network using layer-wise sparse coding and stdp learning. *arXiv [Preprint]*. arXiv:1611.03000.
- Tavanaei, A., and Maida, A. S. (2017). "Multi-layer unsupervised learning in a spiking convolutional neural network," in *2017 International Joint Conference on Neural Networks (IJCNN)* (Anchorage, AK: IEEE), 2023–2030.
- Werbos, P. J. (1990). Backpropagation through time: what it does and how to do it. *Proc. IEEE* 78, 1550–1560.
- Wu, Y., Deng, L., Li, G., Zhu, J., and Shi, L. (2018a). Direct training for spiking neural networks: faster, larger, better. *arXiv [Preprint]*. arXiv:1809.05793.
- Wu, Y., Deng, L., Li, G., Zhu, J., and Shi, L. (2018b). Spatio-temporal backpropagation for training high-performance spiking neural networks. *Front. Neurosci.* 12:331. doi: 10.3389/fnins.2018.00331
- Zhao, B., Ding, R., Chen, S., Linares-Barranco, B., and Tang, H. (2015). Feedforward categorization on aer motion events using cortex-like features in a spiking neural network. *IEEE Trans. Neural Netw. Learn. Syst.* 26, 1963–1978. doi: 10.1109/TNNLS.2014.2362542

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2020 Lee, Sarwar, Panda, Srinivasan and Roy. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.