

Going Deeper With Directly-Trained Larger Spiking Neural Networks

Hanle Zheng¹, Yujie Wu¹, Lei Deng^{1,3}, Yifan Hu¹ and Guoqi Li^{1,2*}

¹ Center for Brain-Inspired Computing Research, Department of Precision Instrument, Tsinghua University, Beijing 100084, China

² Beijing Innovation Center for Future Chip, Tsinghua University, Beijing 100084, China

³ Department of Electrical and Computer Engineering, University of California, Santa Barbara
{zhenghl18, wu-yj16, huyf19}@mails.tsinghua.edu.cn, leideng@ucsb.edu, liguoqi@tsinghua.edu.cn

Abstract

Spiking neural networks (SNNs) are promising in a bio-plausible coding for spatio-temporal information and event-driven signal processing, which is very suited for energy-efficient implementation in neuromorphic hardware. However, the unique working mode of SNNs makes them more difficult to train than traditional networks. Currently, there are two main routes to explore the training of deep SNNs with high performance. The first is to convert a pre-trained ANN model to its SNN version, which usually requires a long coding window for convergence and cannot exploit the spatio-temporal features during training for solving temporal tasks. The other is to directly train SNNs in the spatio-temporal domain. But due to the binary spike activity of the firing function and the problem of gradient vanishing or explosion, current methods are restricted to shallow architectures and thereby difficult in harnessing large-scale datasets (e.g. ImageNet). To this end, we propose a threshold-dependent batch normalization (tdBN) method based on the emerging spatio-temporal backpropagation, termed “STBP-tdBN”, enabling direct training of a very deep SNN and the efficient implementation of its inference on neuromorphic hardware. With the proposed method and elaborated shortcut connection, we significantly extend directly-trained SNNs from a shallow structure (<10 layer) to a very deep structure (50 layers). Furthermore, we theoretically analyze the effectiveness of our method based on “Block Dynamical Isometry” theory. Finally, we report superior accuracy results including 93.15% on CIFAR-10, 67.8% on DVS-CIFAR10, and 67.05% on ImageNet with very few timesteps. To our best knowledge, it’s the first time to explore the directly-trained deep SNNs with high performance on ImageNet. We believe this work shall pave the way of fully exploiting the advantages of SNNs and attract more researchers to contribute in this field.

Introduction

Inspired by human neurons’ working patterns, spiking neural networks (SNNs) have been considered as a promising model in artificial intelligence and theoretical neuroscience (Roy, Jaiswal, and Panda 2019). Benefited from intrinsic neuronal dynamics and event-driven spike communication paradigms, SNNs show great potential in continuous spatio-temporal information processing with lower energy

consumption and better robustness (Stromatiás et al. 2015). Moreover, SNNs can be easily applied on some specialized neuromorphic hardware (Merolla et al. 2014; Davies et al. 2018), which may be seen as the next generation of neural networks.

There are two main approaches to train an SNN with high performance. One is to convert a pre-trained ANN to an SNN model, which usually needs hundreds of timesteps (Sengupta et al. 2019; Hu et al. 2018). So, even though these SNNs achieve comparable accuracy to ANNs with similar structures, the large number of timesteps causes serious signal latency and increase the amount of computation. The other is to train SNNs directly based on gradient descent method, which is independent of pre-trained ANNs and can lessen the number of timesteps. A recent work (Wu et al. 2018), which proposes a learning algorithm called “Spatio-temporal backpropagation” (STBP) with an approach to training SNNs directly on ANN-oriented programming frameworks (e.g. Pytorch), provides us with a chance to explore deeper and larger directly-trained SNNs. However, SNNs trained by this algorithm are restricted to shallow architectures and cannot achieve satisfactory performance on large-scale datasets such as ImageNet. So under above algorithm, we clarify two problems to be solved for training deeper SNNs directly.

The first problem is gradient vanishing or explosion. Because of special mechanism of spatio-temporal information processing and non-differentiable spiking signal, the gradient propagation behaves much unstable and tends to vanish in most cases when training SNNs directly, which prevents SNNs from going deeper. So far, there isn’t an effective method to handle this problem well in directly-trained SNNs. Famous works (Lee et al. 2020; Lee, Delbruck, and Pfeiffer 2016; Wu et al. 2019) before us fail to train deep SNNs directly and all of their models are less than 10 layers, which seriously influences the performances of their methods.

The other problem is that we need to balance the threshold and input on each neuron to get appropriate firing rates in SNNs. When the input is too small compared with the threshold, the neuron fires few spikes and the neuronal membrane potential remains so the information handled by the neuron cannot be expressed enough. When the input is too large, the neuron fires all the time and is insensitive to

*Corresponding author.

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

the change of input. For directly-trained SNNs, with binary spikes propagating layer by layer, the distribution of pre-synaptic inputs will shift during the training process, making the size of inputs inappropriate. Many methods have been proposed to deal with it, such as threshold regularization (Lee, Delbruck, and Pfeiffer 2016) and NeuNorm (Wu et al. 2019).

Normalization seems to be appropriate methods to solve both of the problems. They stabilize the network and gradient propagation according to (Chen et al. 2020). Also, they normalize the distributions of pre-synaptic inputs to same expectation and variance, which helps to balance the threshold and input by reducing the internal covariate shift. However, the existing normalization methods aren't suitable for training of SNNs. For the additional temporal dimension and special activation mechanism, directly-trained SNNs need a specially designed normalization method.

In this paper, we develop a new algorithm to train a deep SNN directly. The main contributions of this work are summarized as follows:

- We propose the threshold-dependent batch normalization to solve the gradient vanishing or explosion problem and adjust firing rate. Furthermore, we take the residual network structure and modify the shortcut connections which are suitable for deep SNNs.
- On this basis, we investigate very deep directly-trained SNNs (**extending them from less than 10 to 50 layers**) and test them over large-scale non-spiking datasets (CIFAR-10, ImageNet) and neuromorphic datasets (DVS-Gesture, DVS-CIFAR10).
- On CIFAR-10 and ImageNet, we comprehensively validate different SNN architectures (ResNet-18, 34, 50) and report competitive results compared to similar SNNs with much fewer timesteps (**no more than 6 timesteps**), to our best knowledge, which is the first time that the directly-trained SNN with full spikes report fairly high accuracy on ImageNet. On neuromorphic datasets, our model achieves state-of-the-art performance on both DVS-Gesture and DVS-CIFAR10, which shows advantage of SNNs on dealing with temporal-spatial information.

Related Work

Learning algorithm of SNNs In the past few years, a lot of learning algorithms have explored how to train a deep SNN with high performance, including: (1) some approaches to converting pre-trained ANNs to SNNs; (2) gradient descent based algorithms.

The first one is called as the "ANN-SNN conversion methods" (Sengupta et al. 2019; Han, Srinivasan, and Roy 2020), seen as the most popular way to train deep SNNs with high performance, which transforms the real-valued output of ReLU function to binary spikes in the SNN. This kind of method successfully reports competitive results over large-scale datasets without serious degradation compared to ANNs. However, it ignores rich temporal dynamic behaviors of spiking neurons and usually requires hundreds or thousands of time steps to approach the accuracy of pre-trained ANNs.

The gradient descent based algorithms train SNNs with error backpropagation. With gradient descent optimization learning algorithms, some SNN models (Lee, Delbruck, and Pfeiffer 2016; Jin, Zhang, and Li 2018; Lee et al. 2020) achieve high performance on CIFAR-10 and other neuromorphic datasets. Among them, (Wu et al. 2019) improves the leaky integrate-and-fire (LIF) model (Hunsberger and Eliasmith 2015) to an iterative LIF model and develop STBP learning algorithm, which makes it friendly on ANN-oriented programming frameworks and speeds up the training process. Moreover, training SNNs directly shows great potential on dealing with spatial and temporal information and reports high accuracy within very few timesteps. However, it fails to train a very deep SNN directly because of the gradient vanishing and internal covariate shift, which is exactly what we want to conquer.

Gradient vanishing or explosion in the deep neural network (DNN) A DNN can avoid gradient vanishing or explosion when it is dynamic isometry, which means every singular value of its input-output jacobian matrix remains close to 1. (Chen et al. 2020) proposes a metric—"Block Dynamical Isometry", serving as a general statistical tool to all of complex serial-parallel DNN. It investigates the first and second moment of each block in the neural network and analyze their effects to the gradient distribution. Moreover, it gives theoretical explanation to the function of the weight initialization, batch normalization and shortcut connection in the neural network, which helps us to develop our algorithm.

Normalization Normalization techniques enable the training of well-behaved neural networks. For artificial neural networks, normalization, such as batch normalization (Ioffe and Szegedy 2015), group normalization (Wu and He 2018), and layer normalization (Ba, Kiros, and Hinton 2016), have become common methods. Batch normalization (BN) accelerates deep networks training by reducing internal covariate shift, which enables higher learning rates and regularizes the model. While it causes high learning latency and increases computation, BN makes it possible for networks to go deeper avoiding gradient vanishing or explosion. For SNNs, researchers propose other normalization techniques, such as data-based normalization (Diehl et al. 2015), Spike-Norm (Sengupta et al. 2019) and NeuNorm (Wu et al. 2019). These normalization methods aim to balance the input and threshold to avoid serious information loss, but they are not effective to our directly-trained deep SNNs because they still neglect the problem of gradient vanishing. We noticed the effects of BN in ANNs and the importance of input distribution in SNNs, so we modify BN to satisfy the training and inference of SNN models.

Materials and Methods

Iterative LIF Model

The iterative LIF model is first proposed by (Wu et al. 2019), who utilizes Euler method to solve the first-order differential

equation of Leaky integrate-and-fire (LIF) model and converts it to an iterative expression

$$u^t = \tau_{decay} u^{t-1} + I^t, \quad (1)$$

where τ_{decay} is a constant to describe how fast the membrane potential decays, u^t is the membrane potential, I^t is the pre-synaptic inputs. Let V_{th} denote the given threshold. When $u^t > V_{th}$, the neuron fires a spike and u^t will be reset to 0. The pre-synaptic inputs are accumulated spikes from other neurons at the last layer. So I^t can be represented by $x^t = \sum_j w_j o^t(j)$, where w_j are weights and $o^t(j)$ denotes binary spiking output from others at the moment of t . Taking spatial structure into consideration and set $u_{reset} = 0$, the whole iterative LIF model in both spatial and temporal domain can be determined by

$$u^{t,n+1} = \tau_{decay} u^{t-1,n+1} (1 - o^{t-1,n+1}) + x^{t,n}, \quad (2)$$

$$o^{t,n+1} = \begin{cases} 1 & \text{if } u^{t,n+1} > V_{th}, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

where $u^{t,n}$ is the membrane potential of the neuron in n -th layer at time t , $o^{t,n}$ is the binary spike and τ_{decay} is the potential decay constant.

The iterative LIF model enables forward and backward propagation to be implemented on both spatial and temporal dimensions, which makes it friendly to general machine-learning programming frameworks.

Threshold-dependent Batch Normalization

As a regular component of DNNs, batch normalization (BN) has been a common method for current neural networks, which allows stable convergence and much deeper neural networks. However, because of the additional temporal dimension and special activation mechanism, directly-trained SNNs need a specially-designed normalization method. That motivates us to propose the threshold-dependent batch normalization (tdBN).

We consider a Spiking Convolution Neural Network (SCNN). Let o^t represent spiking outputs of all neurons in a layer at timestep t . With convolution kernel W and bias B , we have

$$x^t = W \otimes o^t + B, \quad (4)$$

where $x^t \in R^{N \times C \times H \times W}$ represents the pre-synapse inputs at timestep t with N as the batch axis, C as the channel axis, (H, W) as the spatial axis.

In our tdBN, the high-dimensional pre-synaptic inputs will be normalized along the channel dimension (Fig. 1). Let x_k^t represent k -th channel feature maps of x^t . Then $x_k = (x_k^1, x_k^2, \dots, x_k^T)$ will be normalized by

$$\hat{x}_k = \frac{\alpha V_{th} (x_k - E[x_k])}{\sqrt{Var[x_k] + \epsilon}}, \quad (5)$$

$$y_k = \lambda_k \hat{x}_k + \beta_k, \quad (6)$$

where V_{th} denotes the threshold, α is a hyper-parameter depending on network structure, ϵ is a tiny constant, λ_k and β_k are two trainable parameters, $E[x_k]$, $Var[x_k]$ are the mean

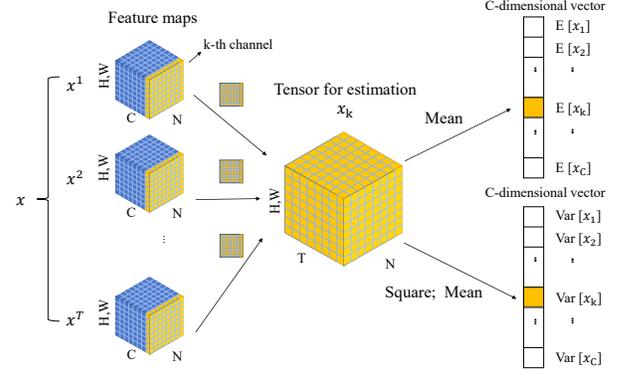


Figure 1: Estimation of $E[x]$ and $Var[x]$ in tdBN. Each cube shows a feature map tensor at t timestep, with N as the batch axis, C as the channel axis, (H, W) as the spatial axis. Each element in C -dimension vector $E[x]$ and $Var[x]$ is estimated by the yellow tensor of corresponding channel.

and variance of x_k statistically estimated over the Mini-Batch. Fig. 1 displays how $E[x_k]$, $Var[x_k]$ to be computed, which are defined by

$$E[x_k] = \text{mean}(x_k), \quad (7)$$

$$Var[x_k] = \text{mean}((x_k - E[x_k])^2). \quad (8)$$

So, during the training, $y_k \in R^{T \times N \times H \times W}$ is exactly the normalized pre-synaptic inputs received by neurons of k -th channel at the next layer during T timesteps.

In the inference, we follow the schema as standard batch normalization to estimate μ_{inf} and σ_{inf}^2 that represent the expectation of $E[x_k]$ and $Var[x_k]$ respectively over the whole datasets, which can be computed during the training process by moving average solution.

Moreover, the batchnorm-scale-fusion is necessary to SNNs with tdBN in inference. It removes the batch normalization operations during the inference, thereby maintaining the network to be full-spiking and enabling it to be implemented on the neuromorphic platforms. Let $W_{c,k}$ and $B_{c,k}$ denote the convolution kernel and bias between the c -th feature map in a layer and the k -th feature map in the next layer. The schema is determined by

$$W'_{c,k} = \lambda_k \frac{\alpha V_{th} W_{c,k}}{\sqrt{\sigma_{inf,k}^2 + \epsilon}}, \quad (9)$$

$$B'_{c,k} = \lambda_k \frac{\alpha V_{th} (B_{c,k} - \mu_{inf,k})}{\sqrt{\sigma_{inf,k}^2 + \epsilon}} + \beta_k, \quad (10)$$

where $W'_{c,k}$ and bias $B'_{c,k}$ denote the transformed weights after the batchnorm-scale-fusion. Thus, during the inference, spikes propagate layer by layer through transformed weights $W'_{c,k}$ and bias $B'_{c,k}$ without batchnorm operations. Therefore, our tdBN only affects the computation costs during training and doesn't influence the running mechanism of SNNs already trained.

In short, our tdBN has two main differences from the standard BN. Firstly, unlike ANNs, SNNs propagate information not only layer by layer but also from last moment to the next. So, tdBN should normalize feature inputs on both temporal and spatial dimensions. Secondly, we make normalized variance dependent on threshold. In tdBN, pre-activations are normalized to $N(0, (\alpha V_{th})^2)$ instead of $N(0, 1)$. And we will initialize the trainable parameters λ and β with 1 and 0. In the serial neural network, the hyper-parameter α is 1. For a local parallel network structure having n branches, α will be $1/\sqrt{n}$. It makes the pre-activations with mean of 0 and standard deviation of V_{th} at the early training. The codes of the tdBN can be found in **Supplementary Material A**.

Overall Training Algorithm

In this section, we present the overall training algorithm of the STBP-tdBN for training deep SNNs from scratch with our tdBN.

In the error backpropagation, we consider the last layer as the decoding layer, and the final output Q can be determined by

$$Q = \frac{1}{T} \sum_{t=1}^T M o^{n,t}, \quad (11)$$

where $o^{n,t}$ is the spike fired by the last output layer, M is the matrix of decoding layer and T denotes the number of timesteps.

Then we make the outputs pass through a softmax layer. The loss function is determined as the cross-entropy. Considering the output $Q = (q_1, q_2, \dots, q_n)$ and label vector $Y = (y_1, y_2, \dots, y_n)$, loss function L is defined by

$$p_i = \frac{e^{q_i}}{\sum_{j=1}^n e^{q_j}}, \quad (12)$$

$$L = - \sum_{i=1}^n y_i \log(p_i). \quad (13)$$

With the iterative LIF model, STBP-tdBN method back-propagates the gradient of the loss L on both spatial and temporal domains. By applying the chain rule, $\frac{\partial L}{\partial o_i^{t,n}}$ and $\frac{\partial L}{\partial u_i^{t,n}}$ can be computed by

$$\frac{\partial L}{\partial o_i^{t,n}} = \sum_{j=1}^{l(n+1)} \frac{\partial L}{\partial u_j^{t,n+1}} \frac{\partial u_j^{t,n+1}}{\partial o_i^{t,n}} + \frac{\partial L}{\partial u_i^{t+1,n}} \frac{\partial u_i^{t+1,n}}{\partial o_i^{t,n}}, \quad (14)$$

$$\frac{\partial L}{\partial u_i^{t,n}} = \frac{\partial L}{\partial o_i^{t,n}} \frac{\partial o_i^{t,n}}{\partial u_i^{t,n}} + \frac{\partial L}{\partial u_i^{t+1,n}} \frac{\partial u_i^{t+1,n}}{\partial u_i^{t,n}}, \quad (15)$$

where $o^{t,n}$ and $u^{t,n}$ represent the spike and membrane potential of the neuron in layer n at time t . Because of the non-differentiable spiking activities, $\frac{\partial o^t}{\partial u^t}$ doesn't exist in fact. To solve this problem, (Wu et al. 2018) proposes derivative curve to approximate the derivative of spiking activity. In this paper, we use the rectangular function, which is proved to be efficient in gradient descent and can be determined by

$$\frac{\partial o^t}{\partial u^t} = \frac{1}{a} \text{sign}(|u^t - V_{th}| < \frac{a}{2}). \quad (16)$$

The codes of the overall training algorithm are shown in **Supplementary Material A**.

Theoretical Analysis

In this section, we will analyze the effects of tdBN to SNNs trained by STBP-tdBN. With theoretical tools about gradient norm theory in ANNs, we find that our tdBN can alleviate the problem of gradient vanishing or explosion during the training process. We also explain the functions of scaling factors α and V_{th} we added during the normalization.

Gradient Norm Theory

The gradient norm theory has been developed well in recent years, which aims to conquer the gradient vanishing or explosion in various neural networks structure. In this paper, we adopt the “**Block Dynamical Isometry**” proposed by (Chen et al. 2020) to analyze tdBN's effect in directly-trained SNNs. It considers a network as a series of blocks

$$f(x) = f_{i,\theta_i} \circ f_{i-1,\theta_{i-1}} \circ \dots \circ f_{1,\theta_1}(x), \quad (17)$$

where the function f_{j,θ_j} represents the j^{th} block and define its input-output jacobian matrix as $\frac{\partial f_j}{\partial f_{j-1}} = J_j$. Let $\phi(J)$ represent the expectation of $\text{tr}(J)$ and $\varphi(J)$ denote $\phi(J^2) - \phi^2(J)$. Then they proved the following lemmas.

Lemma 1. Consider a neural network that can be represented as a series of blocks as Eq. (17) and the j^{th} block's jacobian matrix is denoted as J_j . If $\forall j, \phi(J_j J_j^T) \approx 1$ and $\varphi(J_j J_j^T) \approx 0$, the network achieves “**Block Dynamical Isometry**” and can avoid gradient vanishing or explosion. (Chen et al. 2020)

Lemma 2. Consider a block of neural network, which consists of data normalization with 0-mean, linear transform and rectifier activations (“**General Linear Transform**”). Let 2^{nd} moment of input vector as α_{in} and the output vector as α_{out} , we have $\phi(J J^T) = \frac{\alpha_{out}}{\alpha_{in}}$. (Chen et al. 2020)

Based on the theoretical framework of gradient norm, we combine it with the unique properties of spiking neurons and further analyze the effectiveness of our proposed tdBN algorithm for SNNs.

LIF model has two special hyper-parameters: τ_{decay} and V_{th} , where τ_{decay} influences the gradients propagation in temporal domain and V_{th} effects the spatial dimension. For experiment with SNNs, the τ_{decay} are often set as small value (e.g. 0.25). To analyze the gradient transformation, we simplify the model and set τ_{decay} as zero and we can get the following proposition.

Theorem 1. Consider an SNN with T timesteps and the j^{th} block's jacobian matrix at time t is denoted as J_j^t . When τ_{decay} is equal to 0, if we fix the second moment of input vector and the output vector to V_{th}^2 for each block between two tdBN layers, we have $\phi(J_j^t (J_j^t)^T) \approx 1$ and the training of SNN can avoid gradient vanishing or explosion.

Proof. The proof of **Theorem 1** is based on the **Lemma 1** and **Lemma 2**. The details are presented in **Supplementary Material B**. \square

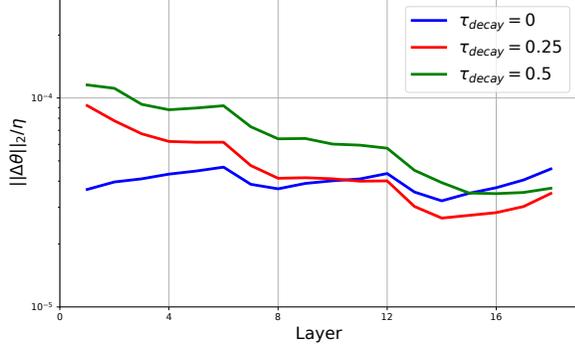


Figure 2: Gradient norm throughout the plain network with tdbN.

Influence of membrane decay mechanism We analyze the effect of τ_{decay} to the gradient propagation. From equations (2) and (15), we have

$$\frac{\partial L}{\partial u_i^{t,n}} = \frac{\partial L}{\partial o_i^{t,n}} \frac{\partial o_i^{t,n}}{\partial u_i^{t,n}} + \frac{\partial L}{\partial u_i^{t+1,n}} \tau_{decay} (1 - o_i^{t,n}). \quad (18)$$

That is to say, if a neuron fires a spike, $(1 - o_i^{t,n})$ is equal to zero and the gradient is irrelevant to τ_{decay} . On the other hand, for τ_{decay} is a tiny constant, the gradient of the neuron at time $t + 1$ has little influence to that at time t .

To verify the **Theorem 1** and our analysis about influence of membrane decay mechanism, we evaluate our tdbN in 20-layers plain spiking network on CIFAR-10. In Fig. 2, we display the mean of gradient norms in each layer during the first 1/6 epoch of training process despite the first encoding layer and the last decoding layer. And we find that when $\tau_{decay} = 0$, the curve of gradient norm behaves much steady, which perfectly supports our theory. It should be emphasized that τ_{decay} cannot be set as 0 because it will prevent the information from propagating along temporal dimension and cause serious degradation. So, we evaluate our method in the condition that $\tau_{decay} \neq 0$. For example, when $\tau_{decay} = 0.25$ and 0.5, the gradient norm increases very slowly as the network deepens, which will not influence the training process. The results strongly support our results that can avoid gradient vanishing or explosion in deep SNNs.

Scaling Factors

As we all know, a key for SNN model to obtain competitive performance is to set suitable threshold to maintain firing rates and reduce information loss. To achieve this, we introduce two scaling factors to the normalization implements in tdbN, which is meant to balance the pre-activations and threshold. In the early training, with two scaling factors— α and V_{th} , we normalize the pre-activations to $N(0, V_{th}^2)$ by initializing the trainable parameters λ and β with 1 and 0.

First, we propose **Theorem 2** to explain the relations between pre-activations and membrane potential of neurons, which helps to understand why our method works.

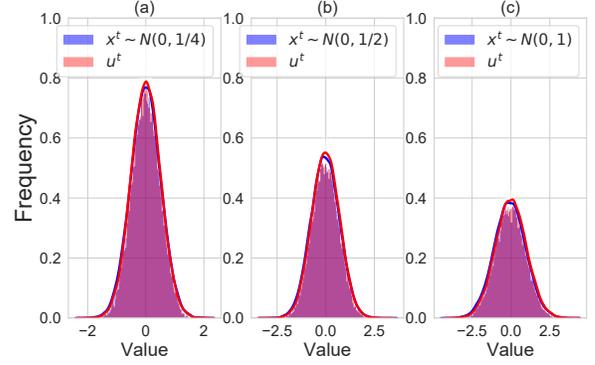


Figure 3: Distributions of membrane potential u^t with different variances of pre-activations x^t when $V_{th} = 1$.

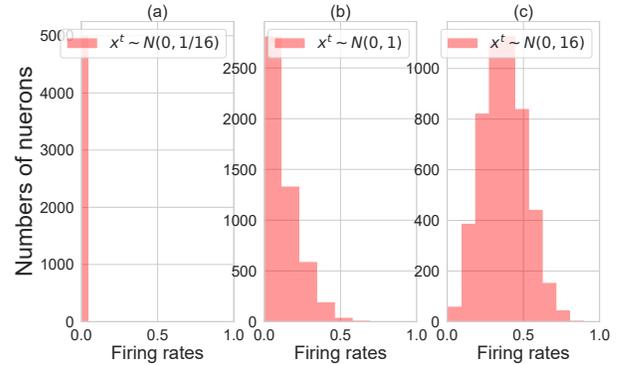


Figure 4: Distributions of neurons' firing rates with different variances of pre-activations x^t when $V_{th} = 1$.

Theorem 2. *With the iterative LIF model, assuming the pre-activations $x^t \sim N(0, \sigma_{in}^2)$, we have the membrane potential $u^t \sim N(0, \sigma_{out}^2)$ and $\sigma_{out}^2 \propto \sigma_{in}^2$.*

Proof. The proof of **Theorem 2** is presented in **Supplementary Material B**. \square

We verify the **Theorem 2** by visualized analysis. In the experiment, we set $\tau_{decay} = 0.25$ and display the distribution of membrane potential with different pre-activations variance σ_{in}^2 . The results are shown in Fig. 3. We find the high degree of similarity between the distributions of pre-activations and membrane potential, which supports the proposition.

Next, we analyze the forward information propagation mechanism with the LIF model.

During the forward, when membrane potential reaches the threshold, neurons will fire a spike and make the information propagate layer by layer. With **Theorem 2** and Eq. (3), we can approximate the possibilities of neurons firing a spike $P(u^t > V_{th})$. It's obvious that it is a positive correlation between $P(u^t > V_{th})$ and the variance of membrane potential σ_{out}^2 as well as σ_{in}^2 . So, we adopt the scaling factors to adjust the distributions of pre-activations to maintain

the firing rates in deep SNNs. Fig. 4 shows the distribution of neurons’ firing rates when we set variances of pre-activations $x^t \sim N(0, \sigma_{in}^2)$ as different values. Because of the decay mechanisms, even a neuron receives positive input each time, it may fire no spike (Fig. 4(a)), which means neurons in next layer only receive few nonzero pre-synaptic inputs, making spikes disappear in deep SNNs and preventing signal from propagating forward. Another situation is that a neuron fires spikes all the time (Fig. 4(c)), which means the outputs of some neurons to be insensitive to the change of pre-activations and causes increase of computation.

In conclusion, to balance pre-synaptic input and threshold to maintain the firing rates, we utilize the scaling factors to control the variance of membrane potential and pre-activations, which alleviates its dependence on the threshold. So, we normalize the pre-activations to $N(0, V_{th}^2)$.

Deep Spiking Residual Network

ResNet is one of the most popular architectures to tackle with the problem of degradation when networks go deep. (He et al. 2016) adds identity mapping between layers, which enables the training of very deep neural networks. Inspired by the residual learning, we propose our deep spiking residual network, which replaces the BN layer with our tdbN and changes the shortcut connection.

Basic block ResNet in ANNs is built with some basic blocks. Fig. 5(a) shows a form of traditional basic blocks in ResNet-ANN. It is implemented in relatively shallow residual networks, which consists of two layers with 3×3 convolution kernel followed by a BN layer and the ReLU activation. On the basis, we propose the basic block of our deep spiking residual network. As shown in Fig. 5(b), we replace the ReLU activation with LIF model and replace the BN layer with our tdbN. Moreover, we change the shortcut connection and add a tdbN layer before the final addition. Hence, expect that the hyper-parameters α in tdbN layers before the final activation layer or in the shortcut are set as $1/\sqrt{2}$, the other tdbN layers’ hyper-parameter α is defined as 1. It guarantees the input distribution of each activation will satisfy $N(0, V_{th}^2)$ at the beginning of training.

Network architecture We build the deep spiking Residual network with our basic blocks. The first encoding layer receives the inputs and performs downsampling directly with convolution kernel and a stride of 2. Then, the spiking activities propagate through basic blocks. Similar to ResNet-ANN, we double the number of channels when the feature map is halved. After average pooling or full-connected layer if necessary, the last decoding layer is a fully-connected layer followed by softmax function.

Experiment

We test our deep residual SNNs on static non-spiking datasets (CIFAR-10, ImageNet) and neuromorphic datasets (DVS-gesture, DVS-CIFAR10). And we compare our results with SNN models to demonstrate the advantages of our method on both accuracy and the number of timesteps.

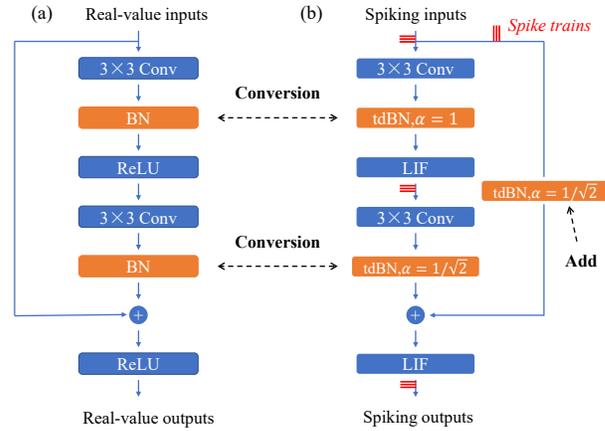


Figure 5: Different basic blocks between traditional ResNet-ANN (a) and our ResNet-SNN with tdbN (b).

The dataset introduction, pre-processing, parameter configuration, training details and results analysis are summarized in **Supplementary Material C**.

Experiment on Static Datasets

Experiments on static datasets include CIFAR-10 and ImageNet, which serve as standard image recognition benchmarks. We test our ResNet-SNN with different timesteps, sizes and depths. With very few timesteps, our models significantly reduce the amount of computation compared with ANNs with same architectures due to sparse spikes. The detailed analysis is presented in **Supplementary Material D**.

CIFAR-10 CIFAR-10 is an image dataset with 50000 training images and 10000 testing images with size of 32×32 , which all belong to 10 classes. In this experiment, we use ResNet-19 with timesteps of 2, 4, 6.

Our results are shown in the Table 1. Before us, SNNs using the ANN-SNN converted method report best accuracy on CIFAR-10. However, to achieve good performance, converted SNNs usually require more than 100 timesteps. In this work, our model achieves state-of-the-art performance (**93.15% top-1 accuracy with only 6 timesteps**) on CIFAR-10, which not only greatly reduces latency and computation cost compared with other SNN models.

ImageNet ImageNet (Deng et al. 2009) contains 1.28 million training images and 50000 validating images. On ImageNet, we test ResNet-34 with standard size and large size. The large model doubles the channels compared with the standard one and achieves **67.05% top-1 accuracy with just 6 timesteps**. Also, we use ResNet-50 to explore the very deep directly-trained SNNs and achieve 64.88% top-1 accuracy. Similar to ResNet-ANN, we observe the enhancement of accuracy between ResNet-34 and ResNet-50, which shows the deeper residual spiking networks may perform better in complex recognition tasks.

Before us, the directly-trained full-spiking SNNs have never reported competitive results on ImageNet. Compared

Dataset	Model	Methods	Architecture	TimeStep	Accuracy
CIFAR-10	(Sengupta et al. 2019)	ANN-SNN	VGG-16	2500	91.55%
	(Hu et al. 2018)	ANN-SNN	ResNet-44	350	92.37%
	(Rathi et al. 2020)	Hybird Training	VGG-16	200	92.02%
	(Lee et al. 2020)	Spike-based BP	ResNet-11	100	90.95%
	(Wu et al. 2019)	STBP	5 Conv, 2Fc	12	90.53%
	our model	STBP-tdBN	ResNet-19	6 4 2	93.16% 92.92% 92.34%
ImageNet	(Sengupta et al. 2019)	ANN-SNN	VGG-16	2500	69.96%
	(Sengupta et al. 2019)	ANN-SNN	ResNet-34	2500	65.47%
	(Han, Srinivasan, and Roy 2020)	ANN-SNN	ResNet-34	1024	66.61%
	(Hu et al. 2018)	ANN-SNN	ResNet-34	768	71.61%
	(Rathi et al. 2020)	Hybird Training	VGG-16	250	65.19%
	(Rathi et al. 2020)	Hybird Training	ResNet-34	250	61.48%
	our model	STBP-tdBN	ResNet-34 ResNet-50 ResNet-34(large)	6 6 6	63.72% 64.88% 67.05%

Table 1: Comparisons with SNNs on CIFAR-10 & ImageNet

with early works, our methods achieve higher accuracies with fewer timesteps. The results are shown in Table 1.

Experiment on Neuromorphic Datasets

Compared with non-spiking static datasets, neuromorphic datasets contain more temporal information, which are more suitable for SNNs to demonstrate their advantages. Previous work (Wu et al. 2019) only tests the MLP-SNN with STBP on N-MNIST. However, N-MNIST is too simple for our deep directly-trained SNNs. So, we adopt two challenging neuromorphic datasets—DVS-Gesture and DVS-CIFAR10.

DVS-Gesture DVS-Gesture (Amir et al. 2017) is a collection of moving gestures performed by 29 different individuals belonging to 11 classes, which is captured by DVS cameras under three lighting conditions. We set timestep T to be 40. In each timestep, the network receives only one slice of the event stream, which means the first 1200ms of each action is used for training or testing. In this experiment, we use ResNet-17 and achieve an accuracy of 96.87%, which is the state-of-the-art result for directly-trained SNNs on DVS-Gesture. We compare our results with other related works on DVS-Gesture, as shown in Table 2.

Model	Methods	Accuracy
(He et al. 2020)	STBP	93.40%
(Shrestha and Orchard 2018)	SLAYER	93.64%
(Kugele et al. 2020)	ANN	95.68%
(Amir et al. 2017)	BPTT	94.59%
our model	STBP-tdBN	96.87%

Table 2: Accuracy comparisons on DVS-Gesture

DVS-CIFAR10 DVS-CIFAR10 (Li et al. 2017) is a neuromorphic dataset converted from famous CIFAR-10 to its dynamic form. It consists of 1000 images in the format of spike train per class. With the noisy environment, it’s also a

challenging recognition task. With ResNet-19, our methods achieve the best performance with 67.8% accuracy. Table 3 shows our results.

Model	Methods	Accuracy
(Sironi et al. 2018)	HATS	52.40%
(Kugele et al. 2020)	Streaming rollout ANN	66.75%
(Ramesh et al. 2019)	DART	65.78%
(Wu et al. 2019)	STBP	60.5%
our model	STBP-tdBN	67.8%

Table 3: Accuracy comparisons on DVS-CIFAR10

Conclusion

In this paper, we present a normalization method enabling directly-trained deep SNNs with high performance. We combine the gradient norm theory and prove that this method can effectively balance the input stimulus and neuronal threshold, thereby facilitating the learning convergence. On this basis, by further introducing the shortcut connection, we greatly extend directly-trained SNNs from a common shallow structure (<10 layer layers) to a very deep structure (> 50 layers). Finally, the model is evaluated on both large-scale static and neuromorphic datasets. Comparing with other SNN models, we achieve a high accuracy on CIFAR-10 and ImageNet with a significantly small inference latency. To our best knowledge, it is the first work to report a directly-trained and very deep SNNs on ImageNet. On neuromorphic datasets, our model can efficiently process temporal-spatial information and achieve state-of-the-art performance.

In summary, this work provides a feasible directly-trained scheme for deep SNNs. It maintains high efficiency of spike-based communication mechanism and enables SNNs to solve more complex recognition tasks, which may benefit the implementations on the neuromorphic hardware and promote the practical applications of SNNs.

Acknowledgments

This work is partially supported by National Key R&D Program of China (No.2018YFE0200200,2018AAA0102600), Beijing Academy of Artificial Intelligence (BAAI), and a grant from the Institute for Guo Qiang of Tsinghua university, and in part by the Science and Technology Major Project of Guangzhou (202007030006), and the key scientific technological innovation research project by Ministry of Education, and the open project of Zhejiang laboratory.

References

- Amir, A.; Taba, B.; Berg, D.; Melano, T.; McKinstry, J.; Di Nolfo, C.; Nayak, T.; Andreopoulos, A.; Garreau, G.; Mendoza, M.; Kusnitz, J.; Debole, M.; Esser, S.; Delbruck, T.; Flickner, M.; and Modha, D. 2017. A Low Power, Fully Event-Based Gesture Recognition System. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Ba, J. L.; Kiros, J. R.; and Hinton, G. E. 2016. Layer Normalization. *arXiv preprint arXiv:1607.06450*.
- Chen, Z.; Deng, L.; Wang, B.; Li, G.; and Xie, Y. 2020. A Comprehensive and Modularized Statistical Framework for Gradient Norm Equality in Deep Neural Networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 1–1.
- Davies, M.; Srinivasa, N.; Lin, T.; Chinya, G.; Cao, Y.; Choday, S. H.; Dimou, G.; Joshi, P.; Imam, N.; Jain, S.; Liao, Y.; Lin, C.; Lines, A.; Liu, R.; Mathaikutty, D.; McCoy, S.; Paul, A.; Tse, J.; Venkataraman, G.; Weng, Y.; Wild, A.; Yang, Y.; and Wang, H. 2018. Loihi: A Neuromorphic Manycore Processor with On-Chip Learning. *IEEE Micro* 38(1): 82–99.
- Deng, J.; Dong, W.; Socher, R.; Li, L.; Kai Li; and Li Fei-Fei. 2009. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 248–255.
- Diehl, P. U.; Neil, D.; Binas, J.; Cook, M.; Liu, S.; 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)*, 1–8.
- Han, B.; Srinivasan, G.; and Roy, K. 2020. RMP-SNN: Residual Membrane Potential Neuron for Enabling Deeper High-Accuracy and Low-Latency Spiking Neural Network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- 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 (CVPR)*.
- He, W.; Wu, Y.; Deng, L.; Li, G.; Wang, H.; Tian, Y.; Ding, W.; Wang, W.; and Xie, Y. 2020. Comparing SNNs and RNNs on Neuromorphic Vision Datasets: Similarities and Differences. *arXiv preprint arXiv:2005.02183*.
- Hu, Y.; Tang, H.; Wang, Y.; and Pan, G. 2018. Spiking Deep Residual Network. *arXiv preprint arXiv:1805.01352*.
- 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*.
- Jin, Y.; Zhang, W.; and Li, P. 2018. Hybrid Macro/Micro Level Backpropagation for Training Deep Spiking Neural Networks. In *Advances in Neural Information Processing Systems 31*, 7005–7015.
- Kugele, A.; Pfeil, T.; Pfeiffer, M.; and Chicca, E. 2020. Efficient Processing of Spatio-Temporal Data Streams With Spiking Neural Networks. *Frontiers in Neuroscience* 14: 439.
- Lee, C.; Sarwar, S. S.; Panda, P.; Srinivasan, G.; and Roy, K. 2020. Enabling Spike-Based Backpropagation for Training Deep Neural Network Architectures. *Frontiers in Neuroscience* 14: 119.
- Lee, J. H.; Delbruck, T.; and Pfeiffer, M. 2016. Training Deep Spiking Neural Networks Using Backpropagation. *Frontiers in Neuroscience* 10: 508.
- Li, H.; Liu, H.; Ji, X.; Li, G.; and Shi, L. 2017. CIFAR10-DVS: An Event-Stream Dataset for Object Classification. *Frontiers in Neuroscience* 11: 309.
- Merolla, P. A.; Arthur, J. V.; Alvarez-Icaza, R.; Cassidy, A. S.; Sawada, J.; Akopyan, F.; Jackson, B. L.; Imam, N.; Guo, C.; Nakamura, Y.; Brezzo, B.; Vo, I.; Esser, S. K.; Appuswamy, R.; Taba, B.; Amir, A.; Flickner, M. D.; Risk, W. P.; Manohar, R.; and Modha, D. S. 2014. A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science* 345(6197): 668–673.
- Ramesh, B.; Yang, H.; Orchard, G. M.; Le Thi, N. A.; Zhang, S.; and Xiang, C. 2019. DART: Distribution Aware Retinal Transform for Event-based Cameras. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 1–1.
- Rathi, N.; Srinivasan, G.; Panda, P.; and Roy, K. 2020. Enabling Deep Spiking Neural Networks with Hybrid Conversion and Spike Timing Dependent Backpropagation. *arXiv preprint arXiv:2005.01807*.
- Roy, K.; Jaiswal, A.; and Panda, P. 2019. Towards spike-based machine intelligence with neuromorphic computing. *Nature* 575(7784): 607–617.
- Sengupta, A.; Ye, Y.; Wang, R.; Liu, C.; and Roy, K. 2019. Going Deeper in Spiking Neural Networks: VGG and Residual Architectures. *Frontiers in Neuroscience* 13: 95.
- Shrestha, S. B.; and Orchard, G. 2018. SLAYER: Spike Layer Error Reassignment in Time. In Bengio, S.; Wallach, H.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 31*, 1412–1421.
- Sironi, A.; Brambilla, M.; Bourdis, N.; Lagorce, X.; and Benosman, R. 2018. HATS: Histograms of Averaged Time Surfaces for Robust Event-Based Object Classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1731–1740.

Stromatias, E.; Neil, D.; Pfeiffer, M.; Galluppi, F.; Furber, S. B.; and Liu, S.-C. 2015. Robustness of spiking Deep Belief Networks to noise and reduced bit precision of neuro-inspired hardware platforms. *Frontiers in Neuroscience* 9: 222.

Wu, Y.; Deng, L.; Li, G.; Zhu, J.; and Shi, L. 2018. Spatio-Temporal Backpropagation for Training High-Performance Spiking Neural Networks. *Frontiers in Neuroscience* 12: 331.

Wu, Y.; Deng, L.; Li, G.; Zhu, J.; Xie, Y.; and Shi, L. 2019. Direct training for spiking neural networks: Faster, larger, better. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 1311–1318.

Wu, Y.; and He, K. 2018. Group normalization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 3–19.