
Going Deeper in Spiking Neural Networks: VGG and Residual Architectures

Abhronil Sengupta¹ Yuting Ye² Robert Wang² Chiao Liu² Kaushik Roy¹

Abstract

Over the past few years, Spiking Neural Networks (SNNs) have become popular as a possible pathway to enable low-power event-driven neuromorphic hardware. However, their application in machine learning have largely been limited to very shallow neural network architectures for simple problems. In this paper, we propose a novel algorithmic technique for generating an SNN with a deep architecture, and demonstrate its effectiveness on complex visual recognition problems such as CIFAR-10 and ImageNet. Our technique applies to both VGG and Residual network architectures, with significantly better accuracy than the state-of-the-art. Finally, we present analysis of the sparse event-driven computations to demonstrate reduced hardware overhead when operating in the spiking domain.

1. Introduction

Spiking Neural Networks (SNNs) are a significant shift from the standard way of operation of Artificial Neural Networks (Farabet et al., 2012). Most of the success of deep learning models of neural networks in complex pattern recognition tasks are based on neural units that receive, process and transmit analog information. Such Analog Neural Networks (ANNs) (Diehl et al., 2015), however, disregard the fact that the biological neurons in the brain (the computing framework after which it is inspired) processes binary spike-based information. Driven by this observation, the past few years have witnessed significant progress in the modeling and formulation of training schemes for SNNs as a new computing paradigm that can potentially replace ANNs as the next generation of Neural Networks. In addition to the fact that SNNs are inherently more biologically plausible, they offer the prospect of event-driven hardware operation. Spiking Neurons process input informa-

tion only on the receipt of incoming binary spike signals. Given a sparsely-distributed input spike train, the hardware overhead (power consumption) for such a spike or event-based hardware would be significantly reduced since large sections of the network that are not driven by incoming spikes can be power-gated (Chen et al., 1998). However, the vast majority of research on SNNs have been limited to very simple and shallow network architectures on relatively simple digit recognition datasets like MNIST (LeCun et al., 1998) while only few works report their performance on more complex standard vision datasets like CIFAR-10 (Krizhevsky & Hinton, 2009) and ImageNet (Russakovsky et al., 2015). The main reason behind their limited performance stems from the fact that SNNs are a significant shift from the operation of ANNs due to their temporal information processing capability. This has necessitated a rethinking of training mechanisms for SNNs.

2. Related Work

Broadly, there are two main categories for training SNNs – supervised and unsupervised. Although unsupervised learning mechanisms like Spike-Timing Dependent Plasticity (STDP) are attractive for the implementation of low-power on-chip local learning, their performance is still outperformed by supervised networks on even simple digit recognition platforms like the MNIST dataset (Diehl & Cook, 2015). Driven by this fact, a particular category of supervised SNN learning algorithms attempts to train ANNs using standard training schemes like backpropagation (to leverage the superior performance of standard training techniques for ANNs) and subsequently convert to event-driven SNNs for network operation (Cao et al., 2015; Zhao et al., 2015; Pérez-Carrasco et al., 2013; Diehl et al., 2015). This can be particularly appealing for NN implementations in low-power neuromorphic hardware specialized for SNNs (Merolla et al., 2014; Akopyan et al., 2015) or interfacing with silicon cochleas or event-driven sensors (Posch et al., 2014; 2011). Our work falls in this category and is based on the ANN-SNN conversion scheme proposed by authors in Ref. (Diehl et al., 2015). However, while prior work considers the ANN operation only during the conversion process, we show that considering the actual SNN operation during the conversion step is crucial for achieving minimal loss in classification accuracy. To

¹Department of Electrical and Computer Engineering, Purdue University, West Lafayette, IN, USA ²Oculus Research, Facebook Research, Redmond, WA, USA. Correspondence to: Abhronil Sengupta <asengup@purdue.edu>.

that effect, we propose a novel weight-normalization technique that ensures that the actual SNN operation is in the loop during the conversion phase. Note that this work tries to exploit neural activation sparsity by converting networks to the spiking domain for power-efficient hardware implementation and are complementary to efforts aimed at exploring sparsity in synaptic connections (Han et al., 2015a).

3. Main Contributions

The specific contributions of our work are as follows:

(i) As will be explained in later sections, there are various architectural constraints involved for training ANNs that can be converted to SNNs in a near-lossless manner. Hence, it is unclear whether the proposed techniques would scale to larger and deeper architectures for more complicated tasks. *We provide, for the first time, proof of concept experiments that deep SNNs (extending from 16 to 34 layers) can provide competitive accuracies over complex datasets like CIFAR-10 and ImageNet.*

(ii) *We propose a new ANN-SNN conversion technique that statistically outperforms state-of-the-art techniques.* We report a classification error of **8.45%** on the CIFAR-10 dataset which is the best-performing result reported for any SNN network, till date. For the first time, we report an SNN performance on the ImageNet 2012 validation set. We achieve a **30.04%** top-1 error rate and **10.99%** top-5 error rate for VGG-16 architectures.

(iii) We explore Residual Network (ResNet) architectures as a potential pathway to enable deeper SNNs. We present insights and design constraints that are required to ensure ANN-SNN conversion for ResNets. We report a classification error of **12.54%** on the CIFAR-10 dataset and a **34.53%** top-1 error rate and **13.67%** top-5 error rate on the ImageNet validation set. *This is the first work that attempts to explore SNNs with residual network architectures.*

(iv) *We demonstrate that SNN network sparsity significantly increases as the network depth increases.* This further motivates the exploration of converting ANNs to SNNs for event-driven operation to reduce compute overhead.

4. Preliminaries

4.1. Input and Output Representation

The main difference between ANN and SNN operation is the notion of time. While ANN inputs are static, SNNs operate based on dynamic binary spiking inputs as a function of time. The neural nodes also receive and transmit binary spike input signals in SNNs, unlike in ANNs, where the inputs and outputs of the neural nodes are analog values. In this work, we consider a rate-encoded network operation

where the average number of spikes transmitted as input to the network over a large enough time window is approximately proportional to the magnitude of the original ANN inputs (pixel intensity in this case). The duration of the time window is dictated by the desired network performance (for instance, classification accuracy) at the output layer of the network. A Poisson event-generation process is used to produce the input spike train to the network. Every time-step of SNN operation is associated with the generation of a random number whose value is compared against the magnitude of the corresponding input. A spike event is triggered if the generated random number is less than the value of the corresponding pixel intensity. This process ensures that the average number of input spikes in the SNN is proportional to the magnitude of the corresponding ANN inputs and is typically used to simulate an SNN for recognition tasks based on datasets for static images (Diehl et al., 2015). Given a Poisson-generated spike train being fed to the network, spikes will be produced at the network outputs. Inference is based on the cumulative spike count of neurons at the output layer of the network over a given time-window.

4.2. ANN and SNN Neural Operation

ANN to SNN conversion schemes usually consider Rectified Linear Unit (ReLU) as the ANN neuron activation function. For a neuron receiving inputs x_i through synaptic weights w_i , the ReLU neuron output y is given by,

$$y = \max \left(0, \sum_i w_i \cdot x_i \right) \quad (1)$$

Although ReLU neurons are typically used in a large number of machine learning tasks at present, the main reason behind their usage for ANN-SNN conversion schemes is that they bear functional equivalence to an Integrate-Fire (IF) Spiking Neuron without any leak and refractory period (Cao et al., 2015; Diehl et al., 2015). Note that this is a particular type of Spiking Neuron model (Izhikevich, 2003). Let us consider the ANN inputs x_i encoded in time as a spike train $\mathbb{X}_i(t)$, where $\mathbb{E}[\mathbb{X}_i(t)] \propto x_i$ (for the rate encoding network being considered in this work). The IF Spiking Neuron keeps track of its membrane potential, v_{mem} , which integrates incoming spikes and generates an output spike whenever the membrane potential cross a particular threshold v_{th} . The membrane potential is reset to zero at the generation of an output spike. All neurons are reset whenever a spike train corresponding to a new image/pattern is presented. The IF Spiking Neuron dynamics as a function of time-step, t , can be described by the following equation,

$$v_{mem}(t+1) = v_{mem}(t) + \sum_i w_i \cdot \mathbb{X}_i(t) \quad (2)$$

Let us first consider the simple case of a neuron being driven by a single input $\mathbb{X}(t)$ and a positive synaptic weight w . Due to the absence of any leak term in the neural dynamics, it is intuitive to show that the corresponding output spiking rate of the neuron is given by $\mathbb{E}[\mathbb{Y}(t)] \propto \mathbb{E}[\mathbb{X}(t)]$, with the proportionality factor being dependent on the ratio of w and v_{th} . In the case when the synaptic weight is negative, the output spiking activity of the IF neuron is zero since the neuron is never able to cross the firing potential v_{th} , mirroring the functionality of a ReLU. The higher the ratio of the threshold with respect to the weight, the more time is required for the neuron to spike, thereby reducing the neuron spiking rate, $\mathbb{E}[\mathbb{Y}(t)]$, or equivalently increasing the time-delay for the neuron to generate a spike. A relatively high firing threshold can cause a huge delay for neurons to generate output spikes. For deep architectures, such a delay can quickly accumulate and cause the network to not produce any spiking outputs for relatively long periods of time. On the other hand, a relatively low threshold causes the SNN to lose any ability to distinguish between different magnitudes of the spike inputs being accumulated to the membrane potential (the term $\sum_i w_i \mathbb{X}_i(t)$ in Eq. 2) of the Spiking Neuron, causing it to lose evidence during the membrane potential integration process. This, in turn, results in accuracy degradation of the converted network. Hence, an appropriate choice of the ratio of the neuron threshold to the synaptic weights is essential to ensure minimal loss in classification accuracy during the ANN-SNN conversion process (Diehl et al., 2015). Consequently, most of the research work in this field has been concentrated on outlining appropriate algorithms for threshold-balancing, or equivalently, weight normalizing different layers of a network to achieve near-lossless ANN-SNN conversion.

4.3. Architectural Constraints

4.3.1. BIAS IN NEURAL UNITS

Typically neural units used for ANN-SNN conversion schemes are trained without any bias term (Diehl et al., 2015). This is due to the fact that optimization of the bias term in addition to the spiking neuron threshold expands the parameter space exploration, thereby causing the ANN-SNN conversion process to be more difficult. Requirement of bias less neural units also entails that Batch Normalization technique (Ioffe & Szegedy, 2015) cannot be used as a regularizer during the training process since it biases the inputs to each layer of the network to ensure each layer is provided with inputs having zero mean. Instead, we use dropout (Srivastava et al., 2014) as the regularization technique. This technique simply masks portions of the input to each layer by utilizing samples from a Bernoulli distribution where each input to the layer has a specified probability of being dropped.

4.3.2. POOLING OPERATION

Deep convolutional neural network architectures typically consist of intermediate pooling layers to reduce the size of the convolution output maps. While various choices exist for performing the pooling mechanism, the two popular choices are either max-pooling (maximum neuron output over the pooling window) or spatial-averaging (two-dimensional average pooling operation over the pooling window). Since the neuron activations are binary in SNNs instead of analog values, performing max-pooling would result in significant information loss for the next layer. Consequently, we consider spatial-averaging as the pooling mechanism in this work (Diehl et al., 2015).

5. Deep Convolutional SNN Architectures: VGG

As mentioned previously, our work is based on the proposal outlined by authors in Ref. (Diehl et al., 2015). In order to ensure that a spiking neuron threshold is sufficiently high to distinguish different magnitude of the spike inputs, a worst case solution would be to set the threshold of a particular layer to the maximum of the summation of all the positive synaptic weights of neurons in that layer. However, such a “Model-Based Normalization” technique is highly pessimistic since all the fan-in neurons are not supposed to fire at every time-step (Diehl et al., 2015). In order to circumvent this issue, authors in Ref. (Diehl et al., 2015) proposed a “Data-Based Normalization” Technique wherein the neuron threshold of a particular layer is set equal to the maximum activation of all ReLUs in the corresponding layer (by passing the entire training set through the trained ANN once after training is completed). Such a “Data-Based” technique performed significantly better than the “Model-Based” algorithm in terms of the final classification accuracy and latency of the converted SNN (three-layered fully connected and convolutional architectures) for a digit recognition problem on the MNIST dataset (Diehl et al., 2015). Note that, this process is referred to as “weight-normalization” and “threshold-balancing” interchangeably in this text. As mentioned before, the goal of this work is to optimize the ratio of the synaptic weights with respect to the neuron firing threshold, v_{th} . Hence, either all the synaptic weights preceding a neural layer are scaled by a normalization factor w_{norm} equal to the maximum neural activation and the threshold is set equal to 1 (“weight-normalization”), or the threshold v_{th} is set equal to the maximum neuron activation for the corresponding layer with the synaptic weights remaining unchanged (“threshold-balancing”). Both operations are exactly equivalent mathematically.

5.1. Proposed Algorithm: SPIKE-NORM

However, the above algorithm leads us to the question: *Are ANN activations representative of SNN activations?* Let us consider a particular example, where the case of maximum activation is shown for a single ReLU. The neuron receives two inputs, namely 0.5 and 1. Let us consider unity synaptic weights in this scenario. Since the maximum ReLU activation is 1.5, the neuron threshold would be set equal to 1.5. However, when this network is converted to the SNN mode, both the inputs would be propagating binary spike signals. The ANN input, equal to 1, would be converted to spikes transmitting at every time-step while the other input would transmit spikes approximately 50% of the duration of a large enough time-window. Hence, the actual summation of spike inputs received by the neuron per time-step would be 2 for a large number of samples, which is higher than the spiking threshold (1.5). Clearly, some information loss would take place due to the lack of this evidence integration.

Driven by this observation, we propose a weight-normalization technique that adaptively balances the threshold of each layer by considering the actual operation of the SNN in the loop during the ANN-SNN conversion process. The algorithm normalizes the weights of the network sequentially for each layer. Given a particular trained ANN, the first step is to generate the input Poisson spike train for the network over the training set for a large enough time-window. The Poisson spike train allows us to record the maximum summation of weighted spike-input (the term $\sum_i w_i \cdot \mathbb{X}_i(t)$ in Eq. 2) that would be received by the first neural layer of the network. In order to minimize the temporal delay of the neuron and simultaneously ensure that the neuron firing threshold is not too low, we weight-normalize the first layer depending on the maximum spike-based input received by the first layer. After the threshold of the first layer is set, we are provided with a representative spike train at the output of the first layer which enables us to generate the input spike-stream for the next layer. The process is continued sequentially for all the layers of the network. The main difference between our proposal and prior work (Diehl et al., 2015) is the fact that the proposed weight-normalization scheme accounts for the actual SNN operation during the conversion process. As we will show in the Results section, this scheme is crucial to ensure near-lossless ANN-SNN conversion for significantly deep architectures and for complex recognition problems.

6. Extension to Residual Architectures

6.1. Architecture Constraints

Residual network architectures were proposed as an attempt to scale convolutional neural networks to very deep

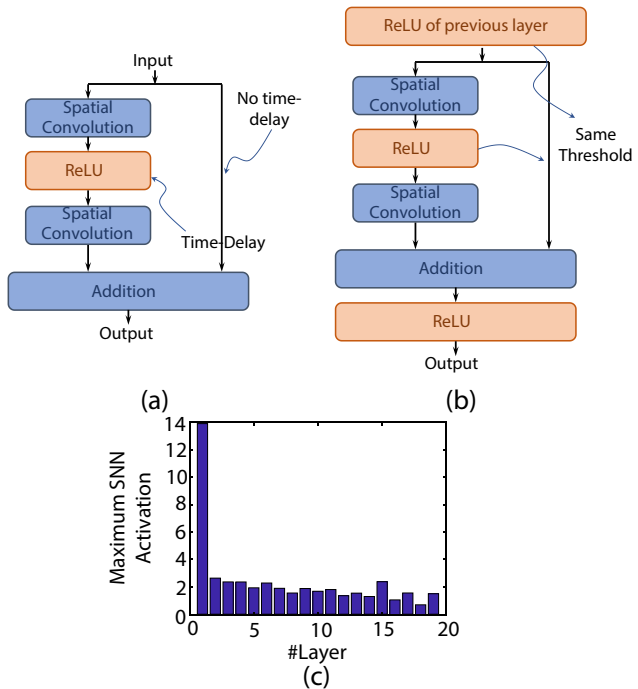


Figure 1. (a) The basic ResNet functional unit, (b) Design constraints introduced in the functional unit to ensure near-lossless ANN-SNN conversion, (c) Typical maximum SNN activations for a ResNet having junction ReLU layers but the non-identity and identity input paths not having the same spiking threshold. While this is not representative of the case with equal thresholds in the two paths, it does justify the claim that after a few initial layers, the maximum SNN activations decay to values close to unity due to the identity mapping.

layered stacks (He et al., 2016a). Although different variants of the basic functional unit have been explored, we will only consider identity shortcut connections in this text (shortcut type-A according to the paper (He et al., 2016a)). Each unit consists of two parallel paths. The non-identity path consists of two spatial convolution layers with an intermediate ReLU layer. While the original ResNet formulation considers ReLUs at the junction of the parallel non-identity and identity paths (He et al., 2016a), recent formulations do not consider junction ReLUs in the network architecture (He et al., 2016b). Absence of ReLUs at the junction point of the non-identity and identity paths was observed to produce a slight improvement in classification accuracy on the CIFAR-10 dataset¹. Due to the presence of the shortcut connections, important design considerations need to be accounted for to ensure near-lossless ANN-SNN conversion. We start with the basic unit, as shown in Fig. 1(a), and point-wise impose various architectural constraints with justifications.

¹<http://torch.ch/blog/2016/02/04/resnets.html>

6.1.1. RELUS AT EACH JUNCTION POINT

As we will show in the Results section, application of our proposed SPIKE-NORM algorithm on such a residual architecture resulted in a converted SNN that exhibited accuracy degradation in comparison to the original trained ANN. We hypothesize that this degradation is attributed mainly to the absence of any ReLUs at the junction points. Each ReLU when converted to an IF Spiking Neuron imposes a particular amount of characteristic temporal delay (time interval between an incoming spike and the outgoing spike due to evidence integration). Due to the shortcut connections, spike information from the initial layers gets instantaneously propagated to later layers. The unbalanced temporal delay in the two parallel paths of the network can result in distortion of the spike information being propagated through the network. Consequently, as shown in Fig. 1(b), we include ReLUs at each junction point to provide a temporal balancing effect to the parallel paths (when converted to IF Spiking Neurons). An ideal solution would be to include a ReLU in the parallel path, but that would destroy the advantage of the identity mapping.

6.1.2. SAME THRESHOLD OF ALL FAN-IN LAYERS

As shown in the next section, direct application of our proposed weight-normalization scheme still resulted in some amount of accuracy loss in comparison to the baseline ANN accuracy. Note that the weight-normalization scheme was applied by considering the immediate fan-in layer, i.e. the junction neuron layer is threshold balanced according to the activity of the non-identity neuron path. However, this layer also receives inputs from the previous junction neuron layer. Since the output spiking activity of a particular neuron is also dependent on the neuron threshold, all the fan-in neuron layers should have the same threshold to ensure that input spike information to the next layer is rate-encoded appropriately. However, the spiking threshold of the neuron layer in the non-identity path is dependent on the activity of the neuron layer at the previous junction. An observation of the typical threshold-balancing factors for the network without using this constraint (shown in Fig. 1(c)) reveal that the threshold-balancing factors mostly lie around unity after a few initial layers. This occurs mainly due to the identity mapping. The maximum summation of spike inputs received by the neurons in the junction layers are dominated by the identity mapping (close to unity). From this observation, we heuristically choose both the thresholds of the non-identity ReLU layer and the identity-ReLU layer equal to 1. However, the accuracy is still unable to approach the baseline ANN accuracy, which leads us to the third design constraint.

6.1.3. INITIAL NON-RESIDUAL PRE-PROCESSING LAYERS

An observation of Fig. 1(c) reveals that the threshold-balancing factors of the initial junction neuron layers are significantly higher than unity. This can be a primary reason for the degradation in classification accuracy of the converted SNN. We note that the residual architectures used by authors in Ref. (He et al., 2016a) use an initial convolution layer with a very wide receptive field (7×7 with a stride of 2) on the ImageNet dataset. The main motive behind such an architecture was to show the impact of increasing depth in their residual architectures on the classification accuracy. Inspired by the VGG-architecture, we replace the first 7×7 convolutional layer by a series of three 3×3 convolutions where the first two layers do not exhibit any shortcut connections. Addition of such initial non-residual pre-processing layers allows us to apply our proposed weight-normalization scheme in the initial layers while using a unity weight-normalization/threshold-balancing factor for the later residual layers. As shown in the Results section, this scheme significantly assists in achieving classification accuracies close to the baseline ANN accuracy since after the initial layers, the maximum neuron activations decay to values close to unity because of the identity mapping.

7. Experiments

7.1. Datasets and Implementation

We evaluate our proposals on standard visual object recognition benchmarks, namely the CIFAR-10 and ImageNet datasets. Experiments performed on networks for the CIFAR-10 dataset are trained on the training set images with per-pixel mean subtracted and evaluated on the testing set. We also present results on the much more complex ImageNet 2012 dataset that contains 1.28 million training images and report evaluation (top-1 and top-5 error rates) on the 50,000 validation set. 224×224 crops from the input images are used for this experiment.

We use VGG-16 architecture (Simonyan & Zisserman, 2014) for both the datasets. ResNet-20 configuration outlined in Ref. (He et al., 2016a) is used for the CIFAR-10 dataset while ResNet-34 is used for experiments on the ImageNet dataset. As mentioned previously, we do not utilize any batch-normalization layers. For VGG networks, a dropout layer is used after every ReLU layer except for those layers which are followed by a pooling layer. For Residual networks, we use dropout only for the ReLUs at the non-identity parallel paths but not at the junction layers. We found this crucial for achieving training convergence.

Our implementation is derived from the Facebook ResNet implementation code for CIFAR and ImageNet datasets

available publicly². We use similar image pre-processing steps and scale and aspect-ratio augmentation techniques as used in (Szegedy et al., 2015). We report single-crop testing results while the error rates can be further reduced with 10-crop testing (Krizhevsky et al., 2012). Networks used for the CIFAR-10 dataset are trained on 2 GPUs with a batchsize of 256 for 200 epochs, while ImageNet training is performed on 8 GPUs for 100 epochs with a similar batchsize. The initial learning rate is 0.05. The learning rate is divided by 10 twice, at 81 and 122 epochs for CIFAR-10 dataset and at 30 and 60 epochs for ImageNet dataset. A weight decay of 0.0001 and a momentum of 0.9 is used for all the experiments. Proper weight initialization is crucial to achieve convergence in such deep networks without batch-normalization. For a non-residual convolutional layer (for both VGG and ResNet architectures) having kernel size $k \times k$ with n output channels, the weights are initialized from a normal distribution and standard deviation $\sqrt{\frac{2}{k^2 n}}$. However, for residual convolutional layers, the standard deviation used for the normal distribution was $\frac{\sqrt{2}}{k^2 n}$. We observed this to be important for achieving training convergence and a similar observation was also outlined in Ref. (Hardt & Ma, 2016) although their networks were trained without both dropout and batch-normalization.

7.2. Experiments for VGG Architectures

Our VGG-16 model architecture follows the implementation outlined in ³ except that we do not utilize the batch-normalization layers. We used a randomly chosen mini-batch of size 256 from the training set for the weight-normalization process on the CIFAR-10 dataset. While the entire training set can be used for the weight-normalization process, using a representative subset did not impact the results. We confirmed this by running multiple independent runs for both the CIFAR and ImageNet datasets. The standard deviation of the final classification error rate after 2500 time-steps was ~ 0.01 . All results reported in this article represent the average of 5 independent runs of the spiking network (since the input to the network is a random process). No notable difference in the classification error rate was observed at the end of 2500 time-steps and the network outputs converged to deterministic values despite being driven by stochastic inputs. For the SNN model based weight-normalization scheme (SPIKE-NORM algorithm) we used 2500 time-steps for each layer sequentially to normalize the weights.

Table 1 summarizes our results for the CIFAR-10 dataset. The baseline ANN error rate on the testing set was 8.3%. Since the main contribution of this work is to minimize the

loss in accuracy during conversion from ANN to SNN for deep-layered networks and not in pushing state-of-the-art results in ANN training, we did not perform any hyperparameter optimization. However, note that despite several architectural constraints being present in our ANN architecture, we are able to train deep networks that provide competitive classification accuracies using the training mechanisms described in the previous subsection. Further reduction in the baseline ANN error rate is possible by appropriately tuning the learning parameters. For the VGG-16 architecture, our implementation of the ANN-model based weight-normalization technique, proposed by Ref. (Diehl et al., 2015), yielded an average SNN error rate of 8.54% leading to an error increment of 0.24%. The error increment was minimized to 0.15% on applying our proposed SPIKE-NORM algorithm. Note that we consider a strict model-based weight-normalization scheme to isolate the impact of considering the effect of an ANN versus our SNN model for threshold-balancing. Further optimizations of considering the maximum synaptic weight during the weight-normalization process (Diehl et al., 2015) is still possible.

Previous works have mainly focused on much shallower convolutional neural network architectures. Although Ref. (Hunsberger & Eliasmith, 2016) reports results with an accuracy loss of 0.18%, their baseline ANN suffers from some amount of accuracy degradation since their networks are trained with noise (in addition to architectural constraints mentioned before) to account for neuronal response variability due to incoming spike trains (Hunsberger & Eliasmith, 2016). It is also unclear whether the training mechanism with noise would scale up to deeper layered networks. *Our work reports the best performance of a Spiking Neural Network on the CIFAR-10 dataset till date. Our network is the deepest convolutional architecture explored so far and our proposed algorithm statistically outperforms state-of-the-art algorithms for ANN-SNN conversion schemes in such architectures.*

The impact of our proposed algorithm is much more apparent on the more complex ImageNet dataset. The rates for the top-1 (top-5) error on the ImageNet validation set are summarized in Table 2. Note that these are single-crop results. The accuracy loss during the ANN-SNN conversion process is minimized by a margin of 0.57% by considering SNN-model based weight-normalization scheme. It is therefore expected that our proposed SPIKE-NORM algorithm would significantly perform better than an ANN-model based conversion scheme as the pattern recognition problem becomes more complex since it accounts for the actual SNN operation during the conversion process. *This is the first work that reports competitive performance of a Spiking Neural Network on the entire 50,000 ImageNet*

²<https://github.com/facebook/fb.resnet.torch>

³<https://github.com/szagoruyko/cifar.torch>

Table 1. Results for CIFAR-10 Dataset

Network Architecture	ANN Error	SNN Error	Error Increment
3-layered networks (Esser et al., 2016)	-	10.68%	-
8-layered networks (Hunsberger & Eliasmith, 2016)	16.28%	16.46%	0.18%
VGG-16 (ANN model based conversion)	8.3%	8.54%	0.24%
VGG-16 (SPIKE-NORM)	8.3%	8.45%	0.15%

2012 validation set. ⁴

7.3. Experiments for Residual Architectures

Our residual networks for CIFAR-10 and ImageNet datasets follow the implementation in Ref. (He et al., 2016a). We first attempt to explain our design choices for ResNets by sequentially imposing each constraint on the network and showing their corresponding impact on network performance in Fig. 2. The “Basic Architecture” involves a residual network without any junction ReLUs. “Constraint 1” involves junction ReLUs without having equal spiking thresholds for all fan-in neural layers. “Constraint 2” imposes an equal threshold of unity for all the layers while “Constraint 3” performs best with two pre-processing plain convolutional layers (3×3) at the beginning of the network. The baseline ANN ResNet-20 was trained with an error of 10.9% on the CIFAR-10 dataset. Note that although we are using terminology consistent with Ref. (He et al., 2016a) for the network architectures, our ResNets contain two extra plain pre-processing layers. The converted SNN according to our proposal yielded a classification error rate of 12.54%. Weight-normalizing the initial two layers using the ANN-model based weight-normalization scheme produced an average error of 12.87%, further validating the efficiency of our weight-normalization technique.

On the ImageNet dataset, we use the deeper ResNet-34 model outlined in Ref. (He et al., 2016a). The initial 7×7 convolutional layer is replaced by three 3×3 convolutional layers where the initial two layers are non-residual plain units. The baseline ANN is trained with an error of 29.31% while the converted SNN error is 34.53% at the end of 2500 timesteps. The results are summarized in Table. 3 and convergence plots for all our networks are provided in Fig. 3.

It is worth noting here that the main motivation of exploring Residual Networks is to go deeper in Spiking Neural Networks. We explore relatively simple ResNet architec-

⁴Ref. (Hunsberger & Eliasmith, 2016) reports a performance of 48.2%(23.8%) on the first 3072-image test batch of the ImageNet 2012 dataset.

Table 2. Results for ImageNet Dataset

Network Architecture	ANN Error	SNN Error	Error Increment
VGG-16 (ANN model based conversion)	29.48% (10.61%)	30.61% (11.21%)	1.13% (0.6%)
VGG-16 (SPIKE-NORM)	29.48% (10.61%)	30.04% (10.99%)	0.56% (0.38%)




Figure 2. Impact of the architectural constraints for Residual Networks. “Basic Architecture” does not involve any junction ReLU layers. “Constraint 1” involves junction ReLUs while “Constraint 2” imposes equal unity threshold for all residual units. Network accuracy is significantly improved with the inclusion of “Constraint 3” that involves pre-processing weight-normalized plain convolutional layers at the network input stage.




Figure 3. Convergence plots for the VGG and ResNet SNN architectures for CIFAR-10 and ImageNet datasets are shown above. The classification error reduces as more evidence is integrated in the Spiking Neurons with increasing time-steps. Note that although the network depths are similar for CIFAR-10 dataset, the ResNet-20 converges much faster than the VGG architecture. The delay for inferencing is higher for ResNet-34 on the ImageNet dataset due to twice the number of layers as the VGG network.

tures, as the ones used in Ref. (He et al., 2016a), which have an order of magnitude lower parameters than standard VGG-architectures. Further hyper-parameter optimizations or more complex architectures are still possible. While the accuracy loss in the ANN-SNN conversion process is more for ResNets than plain convolutional architectures, yet further optimizations like including more pre-processing initial layers or better threshold-balancing schemes for the residual units can still be explored. *This work serves as the first work to explore ANN-SNN conversion schemes for*

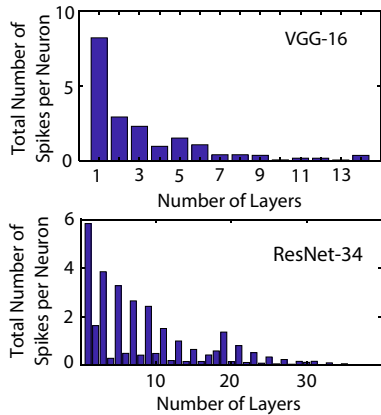


Figure 4. Average cumulative spike count generated by neurons in VGG and ResNet architectures on the ImageNet dataset as a function of the layer number. 500 timesteps were used for accumulating the spike-counts for VGG networks while 2000 time-steps were used for ResNet architectures. The neural spiking sparsity increases significantly as network depth increases.

Residual Networks and attempts to highlight important design constraints required for minimal loss in the conversion process.

7.4. Computation Reduction Due to Sparse Neural Events

ANN operation for prediction of the output class of a particular input requires a single feed-forward pass per image. For SNN operation, the network has to be evaluated over a number of time-steps. However, specialized hardware that accounts for the event-driven neural operation and “computes only when required” can potentially exploit such alternative mechanisms of network operation. For instance, Fig. 4 represents the average total number of output spikes produced by neurons in VGG and ResNet architectures as a function of the layer for ImageNet dataset. A randomly chosen minibatch was used for the averaging process. We used 500 timesteps for accumulating the spike-counts for VGG networks while 2000 time-steps were used for ResNet architectures. This is in accordance to the convergence plots shown in Fig. 3. An important insight obtained from Fig. 4 is the fact that neuron spiking activity becomes sparser as the network depth increases. Hence, benefits from event-driven hardware is expected to increase as the network depth increases. While an estimate of the actual energy consumption reduction for SNN mode of operation is outside the scope of this current work, we provide an intuitive insight by providing the number of computations per synaptic operation being performed in the ANN versus the SNN.

The number of synaptic operations per layer of the network can be easily estimated for an ANN from the architecture

Table 3. Results for Residual Networks

Dataset	Network Architecture	ANN Error	SNN Error
CIFAR-10	ResNet-20	10.9%	12.54%
ImageNet	ResNet-34	29.31% (10.31%)	34.53% (13.67%)

for the convolutional and linear layers. For the ANN, a multiply-accumulate (MAC) computation takes place per synaptic operation. On the other hand, a specialized SNN hardware would perform an accumulate computation (AC) per synaptic operation only upon the receipt of an incoming spike. Hence, the total number of AC operations occurring in the SNN would be represented by the layerwise product and summation of the average cumulative neural spike count for a particular layer and the corresponding number of synaptic operations. Calculation of this metric reveal that for the VGG network the ratio of SNN AC operations to ANN MAC operations is 1.975 while the ratio is 2.4 for the ResNet (the metric includes only ReLU/IF spiking neuron activations in the network). However, note the fact that a MAC operation involves an order of magnitude more energy consumption than an AC operation. For instance, Ref. (Han et al., 2015b) reports that the energy consumption in a 32-bit floating point MAC operation is $3.2pJ$ while the energy consumption is only $0.1pJ$ for an AC operation in 45nm technology. Hence, the energy consumption reduction for our SNN implementation is expected to be $16.2\times$ for the VGG network and $13.3\times$ for the ResNet in comparison to the original ANN implementation.

8. Conclusions and Future Work

This work serves to provide inspiration to the fact that SNNs exhibit similar computing power as their ANN counterparts. This can potentially pave the way for the usage of SNNs in large scale visual recognition tasks, which can be enabled by low-power neuromorphic hardware. However, there are still open areas of exploration for improving SNN performance. A significant contribution to the present success of deep NNs is attributed to Batch-Normalization (Ioffe & Szegedy, 2015). While using bias less neural units constrain us to train networks without Batch-Normalization, algorithmic techniques to implement Spiking Neurons with a bias term should be explored. Further, it is desirable to train ANNs and convert to SNNs without any accuracy loss. Although the proposed conversion technique attempts to minimize the conversion loss to a large extent, yet other variants of neural functionalities apart from ReLU-IF Spiking Neurons could be potentially explored to further reduce this gap. Additionally, further optimizations to minimize the accuracy loss in ANN-SNN conversion for ResNet architectures should be explored to scale SNN performance to even deeper architectures.

References

- Akopyan, Philipp, Sawada, Jun, Cassidy, Andrew, Alvarez-Icaza, Rodrigo, Arthur, John, Merolla, Paul, Imam, Nabil, Nakamura, Yutaka, Datta, Pallab, Nam, Gi-Joon, et al. Truenorth: Design and tool flow of a 65 mw 1 million neuron programmable neurosynaptic chip. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 34(10):1537–1557, 2015.
- Cao, Yongqiang, Chen, Yang, and Khosla, Deepak. Spiking deep convolutional neural networks for energy-efficient object recognition. *International Journal of Computer Vision*, 113(1):54–66, 2015.
- Chen, Zhanping, Johnson, Mark, Wei, Liqiong, and Roy, W. Estimation of standby leakage power in cmos circuit considering accurate modeling of transistor stacks. In *Low Power Electronics and Design, 1998. Proceedings. 1998 International Symposium on*, pp. 239–244. IEEE, 1998.
- Diehl, Peter U and Cook, Matthew. Unsupervised learning of digit recognition using spike-timing-dependent plasticity. *Frontiers in computational neuroscience*, 9, 2015.
- Diehl, Peter U, Neil, Daniel, Binas, Jonathan, Cook, Matthew, Liu, Shih-Chii, and Pfeiffer, Michael. Fast-classifying, high-accuracy spiking deep networks through weight and threshold balancing. In *Neural Networks (IJCNN), 2015 International Joint Conference on*, pp. 1–8. IEEE, 2015.
- Esser, Steven K, Merolla, Paul A, Arthur, John V, Cassidy, Andrew S, Appuswamy, Rathinakumar, Andreopoulos, Alexander, Berg, David J, McKinstry, Jeffrey L, Melano, Timothy, Barch, Davis R, et al. Convolutional networks for fast, energy-efficient neuromorphic computing. *Proceedings of the National Academy of Sciences*, pp. 201604850, 2016.
- Farabet, Clément, Paz, Rafael, Pérez-Carrasco, Jose, Zamarreño-Ramos, Carlos, Linares-Barranco, Alejandro, LeCun, Yann, Culurciello, Eugenio, Serrano-Gotarredona, Teresa, and Linares-Barranco, Bernabe. Comparison between frame-constrained fix-pixel-value and frame-free spiking-dynamic-pixel ConvNets for visual processing. *Frontiers in neuroscience*, 6, 2012.
- Han, Song, Mao, Huizi, and Dally, William J. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149*, 2015a.
- Han, Song, Pool, Jeff, Tran, John, and Dally, William. Learning both weights and connections for efficient neural network. In *Advances in neural information processing systems*, pp. 1135–1143, 2015b.
- Hardt, Moritz and Ma, Tengyu. Identity matters in deep learning. *arXiv preprint arXiv:1611.04231*, 2016.
- He, Kaiming, Zhang, Xiangyu, Ren, Shaoqing, and Sun, Jian. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016a.
- He, Kaiming, Zhang, Xiangyu, Ren, Shaoqing, and Sun, Jian. Identity mappings in deep residual networks. In *European Conference on Computer Vision*, pp. 630–645. Springer, 2016b.
- Hunsberger, Eric and Eliasmith, Chris. Training spiking deep networks for neuromorphic hardware. *arXiv preprint arXiv:1611.05141*, 2016.
- Ioffe, Sergey and Szegedy, Christian. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning*, pp. 448–456, 2015.
- Izhikevich, Eugene M. Simple model of spiking neurons. *IEEE Transactions on neural networks*, 14(6):1569–1572, 2003.
- Krizhevsky, Alex and Hinton, Geoffrey. Learning multiple layers of features from tiny images. 2009.
- Krizhevsky, Alex, Sutskever, Ilya, and Hinton, Geoffrey E. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- LeCun, Yann, Bottou, Léon, Bengio, Yoshua, and Haffner, Patrick. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Merolla, Paul A, Arthur, John V, Alvarez-Icaza, Rodrigo, Cassidy, Andrew S, Sawada, Jun, Akopyan, Philipp, Jackson, Bryan L, Imam, Nabil, Guo, Chen, Nakamura, Yutaka, et al. A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science*, 345(6197):668–673, 2014.
- Pérez-Carrasco, José Antonio, Zhao, Bo, Serrano, Carmen, Acha, Begona, Serrano-Gotarredona, Teresa, Chen, Shouchun, and Linares-Barranco, Bernabé. Mapping from Frame-Driven to Frame-Free Event-Driven Vision Systems by Low-Rate Rate Coding and Coincidence Processing—Application to Feedforward ConvNets. *IEEE transactions on pattern analysis and machine intelligence*, 35(11):2706–2719, 2013.
- Posch, Christoph, Matolin, Daniel, and Wohlgenannt, Rainer. A QVGA 143 dB dynamic range frame-free

PWM image sensor with lossless pixel-level video compression and time-domain CDS. *IEEE Journal of Solid-State Circuits*, 46(1):259–275, 2011.

Posch, Christoph, Serrano-Gotarredona, Teresa, Linares-Barranco, Bernabe, and Delbruck, Tobi. Retinomorphic event-based vision sensors: bioinspired cameras with spiking output. *Proceedings of the IEEE*, 102(10):1470–1484, 2014.

Russakovsky, Olga, Deng, Jia, Su, Hao, Krause, Jonathan, Satheesh, Sanjeev, Ma, Sean, Huang, Zhiheng, Karpathy, Andrej, Khosla, Aditya, Bernstein, Michael, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.

Simonyan, Karen and Zisserman, Andrew. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

Srivastava, Nitish, Hinton, Geoffrey E, Krizhevsky, Alex, Sutskever, Ilya, and Salakhutdinov, Ruslan. Dropout: a simple way to prevent neural networks from overfitting. *Journal of machine learning research*, 15(1):1929–1958, 2014.

Szegedy, Christian, Liu, Wei, Jia, Yangqing, Sermanet, Pierre, Reed, Scott, Anguelov, Dragomir, Erhan, Dumitru, Vanhoucke, Vincent, and Rabinovich, Andrew. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1–9, 2015.

Zhao, Bo, Ding, Ruoxi, Chen, Shoushun, Linares-Barranco, Bernabe, and Tang, Huajin. Feedforward categorization on AER motion events using cortex-like features in a spiking neural network. *IEEE transactions on neural networks and learning systems*, 26(9):1963–1978, 2015.