Neural Architecture Search for Spiking Neural Networks

Youngeun Kim[®], Yuhang Li[®], Hyoungseob Park[®], Yeshwanth Venkatesha, and Priyadarshini Panda[®]

Department of Electrical Engineering Yale University New Haven, CT, USA {youngeun.kim, yuhang.li, hyoungseob.park, yeshwanth.venkatesha, priya.panda}@yale.edu

Abstract. Spiking Neural Networks (SNNs) have gained huge attention as a potential energy-efficient alternative to conventional Artificial Neural Networks (ANNs) due to their inherent high-sparsity activation. However, most prior SNN methods use ANN-like architectures (e.g., VGG-Net or ResNet), which could provide sub-optimal performance for temporal sequence processing of binary information in SNNs. To address this, in this paper, we introduce a novel Neural Architecture Search (NAS) approach for finding better SNN architectures. Inspired by recent NAS approaches that find the optimal architecture from activation patterns at initialization, we select the architecture that can represent diverse spike activation patterns across different data samples without training. Moreover, to further leverage the temporal information among the spikes, we search for feed-forward connections as well as backward connections (*i.e.*, temporal feedback connections) between layers. Interestingly, SNASNet found by our search algorithm achieves higher performance with backward connections, demonstrating the importance of designing SNN architecture for suitably using temporal information. We conduct extensive experiments on three image recognition benchmarks where we show that SNASNet achieves state-of-the-art performance with significantly lower timesteps (5 timesteps). Code is available on Github.

Keywords: Spiking Neural Networks, Neural Architecture Search, Neuromorphic Computing

1 Introduction

Spiking Neural Networks (SNNs) [71, 12, 84, 85, 45, 23] have gained increasing attention as a promising paradigm for low-power intelligence. Inspired by biological neuronal functionality, SNNs process visual information with binary spikes over multiple timesteps. So far, the majority of works on SNNs have focused on image classification problem [71] to develop an energy-efficient alternative to Artificial Neural Networks (ANNs). To this end, recent SNN works utilize ANN architectures (*e.g.*, VGG-Net [77] or ResNet [33]) designed by human experts.



Fig. 1. Accuracy and timesteps for different SNN models on (a) CIFAR10, (b) CIFAR100, and (c) TinyImageNet dataset. While showing comparable accuracy with state-of-the-art networks, SNASNet achieves significantly lower latency. Also, SNASNet-Bw where we search both forward and backward connections provides better performance than the SNASNet-Fw with only forward connections.

While SNNs show an impressive advantage on energy-efficiency, they still lag behind ANNs in terms of accuracy.

In this paper, we argue that the inherent structural/functional difference between ANNs and SNNs induces an unignorable architectural gap, resulting in a sub-optimal solution when we naively deploy ANN architectures on SNNs. Specifically, different from ANNs with ReLU neurons, SNNs consist of Leaky Integrate-and-Fire (LIF) neurons which store and transmit temporal information. However, manually searching for SNN-friendly architectures is laborious. Therefore, we use Neural Architecture Search (NAS) [106, 107, 6, 34, 75, 102], which can automatically discover the optimal SNN architecture. Although NAS has become a prevalent technique in various ANN tasks [11, 9, 26, 101], NAS for designing SNNs has not been investigated. In this work, we ask two questions: Q1. Which NAS algorithm is suitable for SNNs?

Q2. Which SNN architecture provides better performance on an image recognition task?

For the first question, we highlight that the mainstream NAS algorithms either require multiple training stages [106, 107, 2, 105, 78] or require training a supernet once with all architecture candidates [54, 28, 86, 6] which takes longer training time to converge than standard training. As SNNs have a significantly slower training process compared to ANNs (*e.g.*, training SNN with timestep 10 on MNIST with NVIDIA V100 GPU takes $11.43 \times$ more latency compared to the same ANN architecture [52]), the above NAS approaches are difficult to be applied on SNNs. On the other hand, recent works [57, 88, 10] have proposed efficient NAS approaches that search the best neuron cell from initialized networks without any training. Specifically, [57] shows that the network architecture with a high representation power at initialization is likely to achieve higher post-training accuracy. Motivated by this, without the training process, we select the SNN architecture that can represent diverse spike activation patterns across different data samples. To quantify the diversity of networks, we measure the distance of temporal activation patterns between different mini-batch samples. However, SNNs show high sparsity variation on the temporal patterns across different mini-batches, resulting in inaccurate distance measures. To address this, we normalize the distance measure based on the sparsity of given activation patterns, that we term as Sparsity-Aware Hamming Distance (SAHD).

To answer the second question, we search the optimal architecture block for SNNs. Here, we find the connection topology as well as the corresponding operation for each connection following previous works [20, 95]. Different from ANNs, SNNs can leverage backward connections as they convey information through time. The backward connections in SNNs can compute more efficiently because each neuron can participate several times in a network computation [3], and they are likely to capture the temporal correlation of the given input. A line of work has studied backward connections in SNNs with various architectures and training methods [14, 98, 3, 37, 63, 62]. Therefore, we search backward connections as well as forward connections through our NAS algorithm. Surprisingly, SNNs with backward connections yield improved accuracy by up to 3% across various benchmark datasets compared to SNNs with forward connections only. Also, as shown in Fig. 1, *SNASNet* founded by our NAS algorithm achieves state-of-the-art performance with a significantly small number of timesteps.

In summary, our key contributions are as follows: (1) So far, most SNN literature deploys architectures from ANN models which can yield sub-optimal performance for SNNs. For the first time, we showcase a NAS technique for finding better SNN architecture on the image recognition task. (2) Motivated by the prior work [57, 88, 10], we find an SNN-friendly architecture by comparing temporal activation without any training process. Eliminating the training cost to find the optimal architecture brings a huge advantage for SNNs that require significantly longer training time compared to ANNs. (3) We also propose Sparsity-Aware Hamming Distance (SAHD) for addressing sparsity variation of LIF neurons. (4) Furthermore, we search backward connections for leveraging temporal information in spiking inputs, which has not been explored before in NAS approaches for ANN architecture.

2 Related Work

2.1 Spiking Neural Networks

Spiking Neural Networks (SNNs) have gained great attention as an energyefficient alternative over standard Artificial Neural Networks (ANNs) [71, 62, 7, 18, 13, 59, 73, 29, 19, 72, 96, 38, 45, 94, 49, 50, 39, 41, 40, 45, 42, 79, 51, 17]. SNNs process temporal information through weight connections and a Leak-Integrate-and-Fire (LIF) neuron [36] which works as a non-linear activation in SNNs. The LIF neuron has its own memory called membrane potential that can store the temporal spike dynamics by accumulating incoming spike signals. If the membrane potential exceeds a firing threshold, the neuron generates a post-synaptic spike. The integrate-and-fire behavior of neurons induces non-

differentiable transfer function. As a result, standard backpropagation is difficult to be applied during the training phase [61].

To address this, various methods have been proposed to circumvent the nondifferentiable backpropagation problem. Among them, surrogate gradient learning approaches have become popular [48, 47, 61, 74, 27, 84] due to their higher performance and smaller number of timesteps compared to other training techniques. They define a surrogate function for LIF neurons when calculating backward gradients. Wu et al. [85] represent the LIF model in a discrete-time domain and enable SNN training with a Pytorch platform. The authors of [81] propose a training algorithm that calculates backward gradients of the accumulated input and output spikes over the time window. Tandem learning [82, 83] utilizes an auxiliary ANN that facilitates stable error back-propagation for SNN training. A line of work [24, 68] train membrane decay or firing threshold in an LIF neuron. which improves the representation power of SNNs. Also, Batch Norm (BN) [35] has been applied to accelerate the training process of SNNs [46, 39, 104]. In spite of the recent developments in SNN training techniques, all of the prior methods leverage ANN architecture, such as, VGG and ResNet families. We assert that these architectures may provide sub-optimal solution for SNNs. Different from previous methods, we search better SNN architectures for the image recognition task which has not been explored so far. We notify that the concurrent work [60] also aims to find SNN-friendly architecture using evolutionary algorithm. whereas our work is based on NAS without training technique. We also have different search spaces from [60] where we more focus on searching backward connections.

2.2 Neural Architecture Search

Neural Architecture Search (NAS) has been proposed to discover high-performing networks [106, 107, 6, 34, 75, 102]. The early stage of NAS algorithm uses reinforcement learning [106, 107, 2, 105, 78] or evolutionary algorithm [70]. However, such methods require training the searched architecture from scratch for each search step, which is extremely computationally expensive. To address this, weight-sharing approaches have been proposed [6, 54, 80, 86, 4, 5, 66, 100, 28]8, 92]. They train the supernet once which includes all architecture candidates. For instance, Darts [54] jointly optimizes the network parameters and the importance of each architecture candidate. Also, SPOS [28] trains the weight parameters with uniform forward path sampling and finds the optimal architecture via evolutionary strategy. The weight-sharing methods do not require training the architecture from scratch at each search step, resulting in better efficiency compared to previous NAS algorithms. In very recent works, the key focus has been the efficiency of the NAS technique [93, 91, 90, 1, 103] owing to the growing size of dataset and architecture. Interestingly, a line of work suggests the concept of NAS without training where the networks do not require training during the search stage [88, 57, 10]. This can significantly reduce the computational cost for searching optimal architecture. At the same time, several benchmarks [20, 95, 76, 21] have been proposed in order to remove the burden of training

time. Following the success of NAS on image classification domain, NAS has been deployed on various tasks such as object detection [11], segmentation [53, 101], GAN [26], transformer [9], and human pose estimation [97, 89]. Despite the huge progress of NAS algorithm in ANN domain, NAS for SNNs has not been developed yet. In this work, we aim to build better SNN architecture by leveraging NAS. Different from the previous methods that search only forward connections of the networks, we search for backward connections in addition to forward, which furthers leverage the temporal information of spikes.

3 Preliminaries

3.1 Leaky Integrate-and-Fire neuron

Leaky Integrate-and-Fire (LIF) neuron is widely used for constructing SNNs [71, 85, 24]. A neuron has a membrane potential that stores the temporal spike information. We convert the above continuous differential equation into a discrete version as in previous works [85, 24]:

$$u_i^t = (1 - \frac{1}{\tau_m})u_i^{t-1} + \frac{1}{\tau_m}\sum_j w_{ij}o_j^t,$$
(1)

where, u_i^t represents the membrane potential of a neuron *i* at timestep *t*, τ_m is a time constant for decaying the membrane potential. Also, w_{ij} stands for weight connections between neuron *j* and neuron *i*. The neuron *i* accumulates membrane potential and generates a spike output o_i^t whenever membrane potential exceeds the threshold. After firing, the membrane potential is reset to zero.

3.2 NAS without Training

Compared to standard ANNs, SNNs require significantly higher computational cost for training due to multiple feedforward steps [52]. This makes it difficult to search for an optimal SNN architecture with NAS techniques that train the architecture candidate multiple times [106, 107, 2, 105, 78] or train a complex supernet [54, 28, 86, 6]. To minimize the training budget, our work is motivated by the previous works [57, 88, 10] which demonstrate that the optimal architecture can be founded without any training process. Specifically, Mellor *et al.* [57] provide the interesting observation that the architecture having distinctive representations across different data samples is likely to achieve higher post-training performance. To measure the discriminative power of initialized networks, they utilize the activation pattern of ReLU neurons as a binary indicator. If the ReLU neuron generates a positive value (*i.e.*, input > 0), the neuron is mapped to 1; otherwise 0. As a result, ReLU neurons in one layer can be encoded to binary vector **c**. Given N samples in a mini-batch, they construct a kernel matrix by computing Hamming distance $d_H(\mathbf{c}_i, \mathbf{c}_j)$ between different

samples i and j, which can be formulated as follows:

$$\mathbf{K}_{H} = \begin{pmatrix} N_{A} - d_{H}(\mathbf{c}_{1}, \mathbf{c}_{1}) \cdots N_{A} - d_{H}(\mathbf{c}_{1}, \mathbf{c}_{N}) \\ \vdots & \ddots & \vdots \\ N_{A} - d_{H}(\mathbf{c}_{N}, \mathbf{c}_{1}) \cdots N_{A} - d_{H}(\mathbf{c}_{N}, \mathbf{c}_{N}) \end{pmatrix}$$
(2)

Here, N_A stands for the number of ReLU neurons in the given layer. The final score of the architecture candidate is obtained by:

$$s = \log(\det |\sum_{l} \mathbf{K}_{H}^{l}|), \tag{3}$$

where, \mathbf{K}_{H}^{l} is the kernel matrix at layer l. A high score implies low off-diagonal elements of kernel matrix \mathbf{K}_{H} , which means that the activation patterns from different samples are not similar. Finally, the highest-scored architecture among the candidates is selected for training.

4 Methodology

6

In this section, we first introduce a temporal binary indicator of an LIF neuron based on the concept of linear region in neural networks. After that, we present sparsity-aware hamming distance that accounts for the sparsity variation of an LIF neuron. Finally, we provide the search space for our NAS algorithm where we find both forward and backward connections.

4.1 Linear Regions from LIF neurons

NAS without training approaches in ANN domain [57, 10] are based on the theoretical concept of linear region in neural networks [31, 87, 30, 67, 58]. That is, each piecewise linear function (such as, ReLU) divides the input space into multiple linear regions. The composition of multiple piecewise linear functions brings multiple linear regions on the input space. Such a pattern of linear regions is used for measuring the representation power of initialized networks by comparing the patterns between different samples. Here, based on previous work, we introduce the definition of neuron transition (*i.e.*, the boundary of linear region) in a piecewise linear function.

Definition 1. (Raghu et al. [67]) For fixed W, we say a neuron with piecewise linear region transitions between inputs $x, x+\delta$ if its activation function switches linear region between x and $x + \delta$.

For instance, ReLU and Hard Tanh have neuron transition at 0 and $\{-1,1\}$, respectively [67]. Fig. 2(a) also shows the simple example with three ReLU neurons. The input space is divided into two regions by a single ReLU neuron. By composing ReLU neurons, the input space is partitioned into multiple regions where each region represents a different linear function.

According to Definition 1, a LIF neuron can be regarded as a piecewise linear function. For each timestep, the LIF neuron transfers 0 if the membrane



Fig. 2. Illustration of the concept of lin- Fig. 3. Sparsity variance of activation patear regions from ReLU and LIF neurons. Each ReLU (or LIF neuron) divides the two-dimension input space into active and inactive regions.

tern across different layers and different samples in a mini-batch. LIF neuron shows higher variance of activation pattern compared to ReLU neuron.

potential is lower than a firing threshold, otherwise it generates 1 (*i.e.*, spike). Thus, neuron transition occurs when a given input generates an output spike. We illustrate the transfer function of an LIF neuron in Fig. 2(b). Different from ReLU neuron, the output of LIF neuron is not solely dependent on the input. As we shown in Eq. 1, the output of LIF neuron is based on the current input as well as the previous membrane potential. Therefore, the neuron transition point can be changing across time. For example, suppose that the firing threshold is 1 and the membrane potential from the previous timestep is 0.3. In this case, neuron transition happens at input = 0.7. After the neuron fires, the membrane potential is reset to 0, where, the neuron transition point becomes 1. With this time-varying transfer function, the linear region of SNNs becomes more diverse.

Sparsity-Aware Hamming Distance 4.2

In NASWOT [57], Hamming Distance (HD) is a key metric to compare the binary activation pattern $\mathbf{c}_i, \mathbf{c}_j$ between two different mini-batch samples i, j. However, standard HD gives inaccurate distance measurement for SNNs due to the large sparsity variance of binary activation pattern \mathbf{c} of LIF. Here, the term "sparsity" denotes the percentage of 0 in binary activation pattern \mathbf{c} from one layer at a given timestep t. Note, the definition of "sparsity" here is slightly different from the previous works which defines "sparsity" from the activation across all timesteps.

Observation on the sparsity of activation pattern. A ReLU neuron provides a binary activation pattern with about 50% sparsity from Gaussian/Uni-



Fig. 4. Accuracy with respect to architecture score. We randomly select 50 architectures from search space. We show Kendall's τ correlation for quantita- Fig. 5. Illustration of cell-based tive comparison (HD: $\tau = 0.519$, SAHD: $\tau = 0.646$).

neural architecture search.

form weight initialization [32], which is similar across all data samples. On the other hand, a LIF neuron shows a large variation of sparsity across different data samples since the activation pattern is based on the previous membrane potential which is different in each sample. In Fig. 3, we visualize the sparsity of binary activation pattern with 16 mini-batch samples. The results demonstrate that LIF neuron causes a large sparsity variation across different samples.

A problem due to large sparsity variation. This large sparsity variation induces different scales of HD. To explain this, for the data sample i, we model the distribution of LIF neuron output (at each timestep) as i.i.d. Bernoulli distribution where the probability of observing 1 is $1 - r_i^l$:

$$o_i^l \sim \text{Bern}(1 - r_i^l). \tag{4}$$

Here, r_i^l is sparsity of binary activation pattern at layer l. Then, the probability of an activation difference (at the same neuron position) between two data samples i, j can be represented as:

$$Pr(|o_i^l - o_j^l| = 1) = \operatorname{Bern}(r_i^l(1 - r_j^l) + (1 - r_i^l)r_j^l).$$
(5)

Considering that each element of binary activation patterns $\mathbf{c}^l \in \mathbb{R}^{N_A^l}$ is sampled from Bernoulli distribution (Eq. 4), where N_A^l denotes the number of neurons at layer l. Then, the expectation of HD (*i.e.*, $d_H(\mathbf{c}_i^l, \mathbf{c}_i^l)$) can be formulated as:

$$\mathbb{E}[d_H(\mathbf{c}_i^l, \mathbf{c}_j^l)] = N_A^l \mathbb{E}[Pr(|o_i^l - o_j^l| = 1)] = N_A^l \{r_i^l(1 - r_j^l) + (1 - r_i^l)r_j^l\}, \quad (6)$$

Note, all quantities in Eq. 4, Eq. 5, and Eq. 6 are evaluated per timestep, and we average them across timesteps.

As we can observe in Eq. 6, the expectation of HD is the function of sparsity r_i^l and r_i^l . Therefore, HD will provide an inaccurate distance measure for SNN where sparsity r^l has a large variation across data samples (Fig. 3). For example, HD is likely to be small if two activations are in extreme cases, highly-sparse $(r \to 1)$ or highly-dense $(r \to 0)$. On the other hand, HD is likely to be high if two activations are in a moderate range $(r \approx 0.5)$. Thus, based on the sparsity of two activations, HD has a different contribution to the final score s (Eq. 3); the ideal case is when all HD have the same contribution.

The proposed solution. To address this problem, we propose Sparsity-Aware Hamming Distance (SAHD) where Hamming Distance is normalized based on the sparsity of two binary activation patterns. This can be simply done by normalizing the expectation of HD value to a constant α :

$$d_{SAH}(\mathbf{c}_{i}^{l}, \mathbf{c}_{j}^{l}) = \frac{\alpha}{N_{A}^{l} \{r_{i}^{l}(1 - r_{j}^{l}) + (1 - r_{i}^{l})r_{j}^{l}\}} d_{H}(\mathbf{c}_{i}^{l}, \mathbf{c}_{j}^{l}).$$
(7)

We compute the global SAHD score by accumulating layer-wise SAHD across all layers, *i.e.*, $d_{SAH}(\mathbf{c}_i, \mathbf{c}_j) = \sum_l d_{SAH}(\mathbf{c}_i^l, \mathbf{c}_j^l)$. Instead of HD, we use SAHD for computing the kernel matrix (Eq. 2) at each timestep. After that, we sum all kernel matrices to compute the final score using Eq. 3. In Fig. 4, we compare the correlation between architecture score and post-training accuracy for HD and SAHD. The results demonstrate that the proposed SAHD has a higher Kendall's τ value which implies it is a more accurate metric for architecture selection.

4.3 Searching Forward and Backward Connections

Cell-based approach [107, 70, 66, 75, 54, 86] is widely used in NAS research. These methods usually search for the connection topology as well as the corresponding operation for each connection. Then, multiple generated cell architectures construct the whole network. In our search algorithm, we also investigate cell-based architectures. Fig. 5 shows the macro skeleton of our SNN architecture. The first block is the spike encoding layer which directly converts a float value image into spikes like previous works [85, 104, 99]. The main body of the skeleton consists of two searched neuron cells and one reduction cell. The reduction cell includes one convolution layer and 2-by-2 Average pooling with stride 2. Finally, a linear classifier is used for prediction.

Cell Search Strategy. Our cell search space is identical to NAS-Bench-201 [20] (except for backward connections) where each cell includes V = 4 nodes with multiple connections sampled from operation set $O = \{zeroize, skip connection, 1-by-1 convolution, 3-by-3 convolution, 3-by-3 average pooling\}$ (see Fig. 5). Each node contains the sum of all incoming feature maps from edge operation. However, different from [20], we search backward connections in addition to forward connections. In backward operation, we add transformed node feature of *l*-th layer at timestep t - 1 to the node of l'-th (l' < l) layer at timestep t. The backward connections. In Fig. 5, we show examples of cell candidates. In the predefined search space, we select the optimal spiking neuron cell. Reduction cell downsamples the spatial size of the feature map. We do not illustrate Zeroize operation for

Table 1. Classification Accuracy (%) on CIFAR10, CIFAR100, and TinyImageNet.

	Dataset	Training Method	Architecture	Timesteps	Accuracy(%)	
Wu et al. [85]	CIFAR10	Surrogate Gradient	5Conv, 2Linear	12	90.53	
Wu et al. [82]	CIFAR10	Tandem Learning	5Conv, 2Linear	8	89.04	
Rathi et al. [69]	CIFAR10	Hybrid	VGG9	100	90.50	
Han et al. [29]	CIFAR10	ANN-SNN Conversion	VGG16	2048	93.63	
Kundu et al. [44]	CIFAR10	Hybrid	VGG16	100	91.29	
Zheng et al. [104]	CIFAR10	Surrogate Gradient	ResNet19	6	93.16	
Deng <i>et al.</i> [16]	CIFAR10	ANN-SNN Conversion	ResNet20	16	92.42	
Li et al. [49]	CIFAR10	ANN-SNN Conversion	VGG16	32	93.00	
Fang et al. [24]	CIFAR10	Surrogate Gradient	6Conv, 2Linear	8	93.50	
Rathi et al. [68]	CIFAR10	Hybrid	VGG16	5	92.70	
SNASNet-Fw (ours)	CIFAR10	Surrogate Gradient	Searched Architecture	5	93.12 ± 0.42	
SNASNet-Fw (ours)	CIFAR10	Surrogate Gradient	Searched Architecture	8	93.64 ± 0.35	
SNASNet-Bw (ours)	CIFAR10	Surrogate Gradient	Searched Architecture	5	93.73 ± 0.32	
SNASNet-Bw (ours)	CIFAR10	Surrogate Gradient	Searched Architecture	8	94.12 ± 0.25	
Lu and Sengupta [56]	CIFAR100	ANN-SNN Conversion	VGG15	62	63.20	
Park et al. [64]	CIFAR100	TTFS	VGG15	680	68.80	
Rathi et al. [69]	CIFAR100	Hybrid	VGG16	125	67.80	
Han et al. [29]	CIFAR100	ANN-SNN Conversion	VGG16	2048	70.90	
Garg <i>et al.</i> [25]	CIFAR100	DCT	VGG9	48	68.30	
Kundu et al. [44]	CIFAR100	Hybrid	VGG11	120	64.98	
Deng <i>et al.</i> [16]	CIFAR100	ANN-SNN Conversion	ResNet20	32	68.40	
Li et al. [49]	CIFAR100	ANN-SNN Conversion	ResNet20	16	72.33	
Rathi et al. [68]	CIFAR100	Hybrid	VGG16	5	69.67	
SNASNet-Fw (ours)	CIFAR100	Surrogate Gradient	Searched Architecture	5	70.06 ± 0.45	
SNASNet-Bw (ours)	CIFAR100	Surrogate Gradient	Searched Architecture	5	73.04 ± 0.36	
Sengupta et al. [73]	TinyImageNet	ANN-SNN Conversion	VGG11	2500	54.20	
Kundu et al. [44]	TinyImageNet	Hybrid	VGG16	150	51.92	
Garg et al. [25]	TinyImageNet	DCT	VGG13	125	52.43	
SNASNet-Fw (ours)	TinyImageNet	Surrogate Gradient	Searched Architecture	5	52.81 ± 0.56	
SNASNet-Bw (ours)	TinyImageNet	Surrogate Gradient	Searched Architecture	5	54.60 ± 0.48	

simplicity. The forward connections and backward connections can be combined seamlessly. Surprisingly, adding backward connections improves the accuracy of SNNs especially on complex datasets such as CIFAR100 and Tiny-ImageNet. To train the searched SNNs, we use surrogate gradient training [61, 84, 85] (see Supplementary C for details).

5 Experiments

5.1 Implementation Details

Dataset. We evaluate our method on CIFAR10 [43], CIFAR100 [43], TinyImageNet [15]. The details of datasets can be found in Supplementary.

Hyperparameters. Our implementation is based on PyTorch [65]. We train the networks with standard SGD with momentum 0.9, weight decay 0.0005 and also apply random crop and horizontal flip to input images. We set batch size for training as 64. The base learning rate is set to 0.2, 0.1, 0.1 for CIFAR10, CIFAR100, TinyImageNet, respectively. We use cosine learning rate scheduling [55]. Here, we set the total number of epochs to 300, 300, 200, for CIFAR10, CIFAR100, TinyImageNet, respectively. We set τ_m in Eq. 1 to $\frac{4}{3}$. We set α in Eq. 7 to $0.5N_A^l$ to get similar sparsity scale in LIF neuron as a ReLU neuron. Also, we search 5000 architecture candidates from search space (We observe the accuracy saturates after 5000 samples, shown in Supplementary E). We use SpikingJelly [22] package for implementing an LIF neuron.



Fig. 6. Searched architecture examples (forward and backward configuration) for three benchmarks. Blue and red arrows denote forward connection and backward connection, respectively. The number on each arrow represents operations introduced in Fig. 5.

Architectures. Here, we provide details for architectures in Fig. 5. Note, we do not allow two nodes to have both forward and backward connections to ensure training convergence and stability. For the spike encoding layer, we use direct coding [85, 104, 99] where we pass the input image for T time-steps through the first convolution layer which generates spikes. The first neuron cell has C-channel input and C-channel output. Reduction cell consists of Conv(C, 2C)-BN(2C)-LIF followed by AvgPool(2). The second neuron cell has 2C-channel input and 2C-channel output. Note, the structures of the first neuron cell and second neuron cell are identical. We set C to 256, 128, 128 for CIFAR10, CIFAR100, TinyI-magNet, respectively. For vectorize block, we first apply AvgPool(2) to the input feature and vectorize the output. Finally, the classifier consists of Dropout(0.5)-FC(1024)-Voting layer, where a voting layer is used to improve the robustness of classification [24].

5.2 Performance Comparison

Table 1 shows the performance comparison between our SNASNet founded by the proposed NAS algorithm and previous SNN models on three benchmarks. As our NAS approach has randomness, we run the same configuration 5 times and report the mean and standard deviation. In the table, "SNASNet-Fw" refers to our searched model with only forward connections and "SNASNet-Bw" denotes our searched model with both forward and backward connections. SNASNet-Fw achieves comparable performance with the previous works with extremely small timesteps. For example, our searched model achieves 70.06% with timestep 5 on CIFAR100, which is similar to the VGG16 model performance from Rathi et al. [68]. Note that, for CIFAR10, which is a relatively simple dataset, a few methods yield marginally better performance than SNASNet-Fw. Interestingly, compared to SNASNet-Fw, SNASNet-Bw improves the performance by 0.61%, 2.98%, and 1.79%, for CIFAR10, CIFAR100, and TinyImageNet, respectively. We note that SNASNet-Bw yields SOTA results across all datasets with only 5/8 timesteps. The results support our assertion that the representation power of SNNs can be enhanced by passing information through backward connections where temporal information is further exploited. We also illustrate the example of



Fig. 7. Accuracy vs. cell attributes. 1st & 2nd columns: Searching Fw connection only. 3rd~7th columns: Searching both Fw and Bw connections. 1st, 2nd, 3rd rows show statistics of Fw, Bw, Fw+Bw, respectively.

searched architecture cell found by our proposed NAS algorithm for each dataset in Fig. 6. Recently, Shu *et al.* [75] show that fast convergence ANN architectures bring smooth loss landscape and accurate gradient information, resulting in high test accuracy. We also found that our searched SNN architectures achieve fast convergence with high test accuracy, as shown in Fig. 8. By using this early stage information, there is a possibility of applying an evolutionary algorithm [70] to SNN searching in future works.

5.3 Experimental Analysis

Observations from Searched Cells. We provide several observations obtained from our searching algorithm. To this end, in Fig. 7, we ran 100 random searches on CIFAR100 and provide averaged accuracy with respect to the number of forward connections (Fw), backward connections (Bw), skip connections, Conv 3×3 , Conv 1×1 , Average pooling. The key observations are as follows. For **SNASNet with only forward connections**, (1) a deeper and wider cell improves performance, which implies that scaling up SNN is important (1st & 2nd columns in Fig. 7). (2) Convolutional layers are important for getting higher performance. On the other hand, average pooling is not preferred for SNNs. For SNASNet with both forward and backward connections (3rd~7th columns in Fig. 7), (1) The trend of forward connections (1st row) also prefer convolutional layers, which is similar to that of 1st row in Fig. 7. (2) As shown in the 2nd row, a small number of backward connections are preferred. Also, the type of connections does not affect the accuracy except for skip connection. (3) More than 2 backward skip connections degrade accuracy significantly. This implies that feedback without transformation (e.g., convolutional or pooling op-

Celltype	Searching dataset	Train/test dataset	Accuracy (%)	$\Delta Acc \ (\%)$
Forward	CIFAR10	CIFAR100	69.98	-0.22
Forward	CIFAR10	TinyImageNet	52.02	-0.56
Forward	CIFAR100	CIFAR10	93.12	+0.01
Forward	CIFAR100	TinyImageNet	52.28	-0.30
Forward	TinyImageNet	CIFAR10	93.16	+0.05
Forward	TinyImageNet	CIFAR100	70.31	+0.11
Backward	CIFAR10	CIFAR100	73.19	-0.12
Backward	CIFAR10	TinyImageNet	54.61	-0.80
Backward	CIFAR100	CIFAR10	93.73	-0.09
Backward	CIFAR100	TinyImageNet	56.00	+0.59
Backward	TinyImageNet	CIFAR10	93.64	-0.18
Backward	TinyImageNet	CIFAR100	73.14	-0.17

Table 2. Transferability study of founded architectures. Δ Acc denotes the performance change compared to the best performed architecture shown in Fig. 6.



Fig. 8. Comparison of (a) training loss and (b) training accuracy across different searched architectures on CIFAR100.

Fig. 9. (a) Accuracy and (b) search time with respect to number of search samples. We run the same settings 5 times.

eration) deteriorates representation of SNNs. (4) Considering both forward and backward connections (3rd row), the total number of connections should be carefully designed.

Transferability of Searched Architecture. We conduct transferability analysis on searched SNN architecture in order to check the dependency of our searching method on the dataset. We search the optimal architecture on dataset A and train/test the searched architecture on dataset B ($A \neq B$). In Table 2, for both forward and backward configurations, we use the searched architecture (Fig. 6) for all experiments. Note, ΔAcc in 2 denotes the performance difference between the accuracy of the original searched architecture and transferred architecture. The results show that the searched SNASNets are surprisingly transferable across diverse datasets, which opens up the promising advantage of eliminating searching time for huge and complex datasets.

Analysis on Distance Metric. In our method, we use we propose Sparsity-Aware Hamming Distance (SAHD) where Hamming Distance (HD) is normalized based on the sparsity of the given two binary activation patterns. This effectively addresses a large variation of spike activation across data samples. We evaluate the performance of SNASNet according to the distance metric: HD **Table 3.** Performance comparison between HD and SAHD on CIFAR100.

Table 4. Ablation studies on timestepson CIFAR100.

Architecture	HD	SAHD	1	Architecture	T=5	T = 10	T=15	T=20
SNASNet-Fw SNASNet-Bw	$\begin{array}{c} 64.16 \pm 2.02 \\ 66.80 \pm 1.73 \end{array}$	$\begin{array}{c} 70.06 \pm 0.45 \\ 73.04 \pm 0.36 \end{array}$		SNASNet-Fw SNASNet-Bw	$70.06 \\ 73.04$	$70.08 \\ 73.46$	$70.56 \\ 73.49$	$70.52 \\ 74.24$

vs. SAHD. In Table 3, we report the performance of SNASNet-Fw and SNASNet-Bw on CIFAR100. Here, we also run the search algorithm 5 times and report the mean and standard deviation. The results demonstrate that SAHD reveals better architecture with less standard deviation in terms of test accuracy for both SNASNet-Fw and SNASNet-Bw architectures.

Ablation on Number of Search Samples. In Fig. 9(a), we report the accuracy with respect to the number of search samples used in our searching algorithm. We change the number of search samples [50, 100, 500, 1000, 5000] on CIFAR100. The backward connection configuration (marked as red) shows higher variation as well as higher performance increase compared to that of the forward connection setting. This is because searching backward connections has larger search space than searching forward connections only. We also measure the computational time for searching in Fig. 9(b). We conduct the experiments on NVIDIA RTX 2080ti GPU and Intel(R) Xeon(R) Gold 6240 CPU @ 2.60GHz processor. The results show that searching backward connection requires sightly longer time than searching forward connection.

Analysis on Timesteps. In Table 4, we report the performance on CIFAR100 with respect to the number of timesteps used in SNNs. Both SNASNet-Fw and SNASNet-Bw achieve performance gain with more number of timesteps. SNASNet-Fw and SNASNet-Bw with 20 timesteps have improved accuracy by 0.48% and 1.2% compared to 5 timesteps, respectively. Interestingly, the performance gain from SNASNet-Bw is larger than SNASNet-Fw. The results suggest that adding backward connections to SNNs effectively leverages the temporal information for improved learning, and thus supports the advantage of backward connections in SNNs.

6 Conclusion

In this paper, we search better SNN architecture using the temporal activation pattern of initialized network. Our search space considers backward search connections in addition to forward connections, which brings the benefit of using temporal information. By achieving better performance than the previous works, we demonstrate that a new type of architecture is more suitable for SNNs where spikes convey information through multiple timesteps.

Acknowledgment. This work was supported in part by C-BRIC, a JUMP center sponsored by DARPA and SRC, Google Research Scholar Award, the National Science Foundation (Grant#1947826), TII (Abu Dhabi) and the DARPA AI Exploration (AIE) program.

References

- Abdelfattah, M.S., Mehrotra, A., Dudziak, L., Lane, N.D.: Zero-cost proxies for lightweight nas. arXiv preprint arXiv:2101.08134 (2021)
- Baker, B., Gupta, O., Naik, N., Raskar, R.: Designing neural network architectures using reinforcement learning. arXiv preprint arXiv:1611.02167 (2016)
- Bellec, G., Scherr, F., Subramoney, A., Hajek, E., Salaj, D., Legenstein, R., Maass, W.: A solution to the learning dilemma for recurrent networks of spiking neurons. Nature communications 11(1), 1–15 (2020)
- Bender, G., Kindermans, P.J., Zoph, B., Vasudevan, V., Le, Q.: Understanding and simplifying one-shot architecture search. In: International Conference on Machine Learning. pp. 550–559. PMLR (2018)
- 5. Brock, A., Lim, T., Ritchie, J.M., Weston, N.: Smash: one-shot model architecture search through hypernetworks. arXiv preprint arXiv:1708.05344 (2017)
- Cai, H., Zhu, L., Han, S.: Proxylessnas: Direct neural architecture search on target task and hardware. arXiv preprint arXiv:1812.00332 (2018)
- Cao, Y., Chen, Y., Khosla, D.: Spiking deep convolutional neural networks for energy-efficient object recognition. International Journal of Computer Vision 113(1), 54–66 (2015)
- Chen, B., Li, P., Li, B., Lin, C., Li, C., Sun, M., Yan, J., Ouyang, W.: Bnnas: Neural architecture search with batch normalization. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 307–316 (2021)
- Chen, B., Li, P., Li, C., Li, B., Bai, L., Lin, C., Sun, M., Yan, J., Ouyang, W.: Glit: Neural architecture search for global and local image transformer. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 12–21 (2021)
- Chen, W., Gong, X., Wang, Z.: Neural architecture search on imagenet in four gpu hours: A theoretically inspired perspective. arXiv preprint arXiv:2102.11535 (2021)
- Chen, Y., Yang, T., Zhang, X., Meng, G., Xiao, X., Sun, J.: Detnas: Backbone search for object detection. Advances in Neural Information Processing Systems 32, 6642–6652 (2019)
- Christensen, D.V., Dittmann, R., Linares-Barranco, B., Sebastian, A., Le Gallo, M., Redaelli, A., Slesazeck, S., Mikolajick, T., Spiga, S., Menzel, S., et al.: 2022 roadmap on neuromorphic computing and engineering. Neuromorphic Computing and Engineering (2022)
- Comsa, I.M., Fischbacher, T., Potempa, K., Gesmundo, A., Versari, L., Alakuijala, J.: Temporal coding in spiking neural networks with alpha synaptic function. In: ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 8529–8533. IEEE (2020)
- Demin, V., Nekhaev, D.: Recurrent spiking neural network learning based on a competitive maximization of neuronal activity. Frontiers in neuroinformatics 12, 79 (2018)
- Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A largescale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
- Deng, S., Gu, S.: Optimal conversion of conventional artificial neural networks to spiking neural networks. arXiv preprint arXiv:2103.00476 (2021)
- 17. Deng, S., Li, Y., Zhang, S., Gu, S.: Temporal efficient training of spiking neural network via gradient re-weighting. arXiv preprint arXiv:2202.11946 (2022)

- 16 Y. Kim, Y. Li, H. Park, V. Yeshwanth, P. Panda.
- Diehl, P.U., Cook, M.: Unsupervised learning of digit recognition using spiketiming-dependent plasticity. Frontiers in computational neuroscience 9, 99 (2015)
- Diehl, P.U., Neil, D., Binas, J., Cook, M., Liu, S.C., Pfeiffer, M.: Fast-classifying, high-accuracy spiking deep networks through weight and threshold balancing. In: 2015 International Joint Conference on Neural Networks (IJCNN). pp. 1–8. ieee (2015)
- 20. Dong, X., Yang, Y.: Nas-bench-201: Extending the scope of reproducible neural architecture search. arXiv preprint arXiv:2001.00326 (2020)
- Duan, Y., Chen, X., Xu, H., Chen, Z., Liang, X., Zhang, T., Li, Z.: Transnasbench-101: Improving transferability and generalizability of cross-task neural architecture search. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5251–5260 (2021)
- 22. Fang, W., Chen, Y., Ding, J., Chen, D., Yu, Z., Zhou, H., Tian, Y., other contributors: Spikingjelly. https://github.com/fangwei123456/spikingjelly (2020)
- Fang, W., Yu, Z., Chen, Y., Huang, T., Masquelier, T., Tian, Y.: Deep residual learning in spiking neural networks. arXiv preprint arXiv:2102.04159 (2021)
- Fang, W., Yu, Z., Chen, Y., Masquelier, T., Huang, T., Tian, Y.: Incorporating learnable membrane time constant to enhance learning of spiking neural networks. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2661–2671 (2021)
- Garg, I., Chowdhury, S.S., Roy, K.: Dct-snn: Using dct to distribute spatial information over time for low-latency spiking neural networks. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4671–4680 (2021)
- Gong, X., Chang, S., Jiang, Y., Wang, Z.: Autogan: Neural architecture search for generative adversarial networks. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3224–3234 (2019)
- Gu, P., Xiao, R., Pan, G., Tang, H.: Stca: Spatio-temporal credit assignment with delayed feedback in deep spiking neural networks. In: IJCAI. pp. 1366–1372 (2019)
- Guo, Z., Zhang, X., Mu, H., Heng, W., Liu, Z., Wei, Y., Sun, J.: Single path oneshot neural architecture search with uniform sampling. In: European Conference on Computer Vision. pp. 544–560. Springer (2020)
- Han, B., Srinivasan, G., Roy, K.: Rmp-snn: Residual membrane potential neuron for enabling deeper high-accuracy and low-latency spiking neural network. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13558–13567 (2020)
- Hanin, B., Rolnick, D.: Complexity of linear regions in deep networks. In: International Conference on Machine Learning. pp. 2596–2604. PMLR (2019)
- Hanin, B., Rolnick, D.: Deep relu networks have surprisingly few activation patterns (2019)
- He, K., Zhang, X., Ren, S., Sun, J.: Delving deep into rectifiers: Surpassing humanlevel performance on imagenet classification. In: Proceedings of the IEEE international conference on computer vision. pp. 1026–1034 (2015)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778 (2016)
- 34. Hu, S., Xie, S., Zheng, H., Liu, C., Shi, J., Liu, X., Lin, D.: Dsnas: Direct neural architecture search without parameter retraining. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12084–12092 (2020)
- 35. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. arXiv preprint arXiv:1502.03167 (2015)

- Izhikevich, E.M.: Simple model of spiking neurons. IEEE Transactions on neural networks 14(6), 1569–1572 (2003)
- Jia, S., Zhang, T., Cheng, X., Liu, H., Xu, B.: Neuronal-plasticity and rewardpropagation improved recurrent spiking neural networks. Frontiers in Neuroscience 15, 205 (2021)
- Jin, X., Rast, A., Galluppi, F., Davies, S., Furber, S.: Implementing spike-timingdependent plasticity on spinnaker neuromorphic hardware. In: The 2010 International Joint Conference on Neural Networks (IJCNN). pp. 1–8. IEEE (2010)
- 39. Kim, Y., Panda, P.: Revisiting batch normalization for training low-latency deep spiking neural networks from scratch. arXiv preprint arXiv:2010.01729 (2020)
- 40. Kim, Y., Panda, P.: Optimizing deeper spiking neural networks for dynamic vision sensing. Neural Networks (2021)
- Kim, Y., Panda, P.: Visual explanations from spiking neural networks using interspike intervals. Sci Rep 11, 19037 (2021). https://doi.org/10.1038/s41598-021-98448-0 (2021)
- Kim, Y., Venkatesha, Y., Panda, P.: Privatesnn: Fully privacy-preserving spiking neural networks. arXiv preprint arXiv:2104.03414 (2021)
- Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009)
- 44. Kundu, S., Datta, G., Pedram, M., Beerel, P.A.: Spike-thrift: Towards energyefficient deep spiking neural networks by limiting spiking activity via attentionguided compression. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 3953–3962 (2021)
- 45. Kundu, S., Pedram, M., Beerel, P.A.: Hire-snn: Harnessing the inherent robustness of energy-efficient deep spiking neural networks by training with crafted input noise. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5209–5218 (2021)
- Ledinauskas, E., Ruseckas, J., Juršėnas, A., Buračas, G.: Training deep spiking neural networks. arXiv preprint arXiv:2006.04436 (2020)
- Lee, C., Sarwar, S.S., Panda, P., Srinivasan, G., Roy, K.: Enabling spike-based backpropagation for training deep neural network architectures. Frontiers in Neuroscience 14 (2020)
- Lee, J.H., Delbruck, T., Pfeiffer, M.: Training deep spiking neural networks using backpropagation. Frontiers in neuroscience 10, 508 (2016)
- Li, Y., Deng, S., Dong, X., Gong, R., Gu, S.: A free lunch from ann: Towards efficient, accurate spiking neural networks calibration. arXiv preprint arXiv:2106.06984 (2021)
- Li, Y., Deng, S., Dong, X., Gu, S.: Converting artificial neural networks to spiking neural networks via parameter calibration. arXiv preprint arXiv:2205.10121 (2022)
- Li, Y., Guo, Y., Zhang, S., Deng, S., Hai, Y., Gu, S.: Differentiable spike: Rethinking gradient-descent for training spiking neural networks. Advances in Neural Information Processing Systems 34, 23426–23439 (2021)
- 52. Liang, L., Qu, Z., Chen, Z., Tu, F., Wu, Y., Deng, L., Li, G., Li, P., Xie, Y.: H2learn: High-efficiency learning accelerator for high-accuracy spiking neural networks. arXiv preprint arXiv:2107.11746 (2021)
- Liu, C., Chen, L.C., Schroff, F., Adam, H., Hua, W., Yuille, A.L., Fei-Fei, L.: Autodeeplab: Hierarchical neural architecture search for semantic image segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 82–92 (2019)

- 18 Y. Kim, Y. Li, H. Park, V. Yeshwanth, P. Panda.
- Liu, H., Simonyan, K., Yang, Y.: Darts: Differentiable architecture search. arXiv preprint arXiv:1806.09055 (2018)
- 55. Loshchilov, I., Hutter, F.: Sgdr: Stochastic gradient descent with warm restarts. arXiv preprint arXiv:1608.03983 (2016)
- Lu, S., Sengupta, A.: Exploring the connection between binary and spiking neural networks. Frontiers in Neuroscience 14, 535 (2020)
- Mellor, J., Turner, J., Storkey, A., Crowley, E.J.: Neural architecture search without training. In: International Conference on Machine Learning. pp. 7588–7598. PMLR (2021)
- 58. Montúfar, G., Pascanu, R., Cho, K., Bengio, Y.: On the number of linear regions of deep neural networks. arXiv preprint arXiv:1402.1869 (2014)
- Mostafa, H.: Supervised learning based on temporal coding in spiking neural networks. IEEE transactions on neural networks and learning systems 29(7), 3227– 3235 (2017)
- 60. Na, B., Mok, J., Park, S., Lee, D., Choe, H., Yoon, S.: Autosnn: Towards energyefficient spiking neural networks. arXiv preprint arXiv:2201.12738 (2022)
- Neftci, E.O., Mostafa, H., Zenke, F.: Surrogate gradient learning in spiking neural networks. IEEE Signal Processing Magazine 36, 61–63 (2019)
- 62. Panda, P., Aketi, S.A., Roy, K.: Toward scalable, efficient, and accurate deep spiking neural networks with backward residual connections, stochastic softmax, and hybridization. Frontiers in Neuroscience 14 (2020)
- Panda, P., Roy, K.: Learning to generate sequences with combination of hebbian and non-hebbian plasticity in recurrent spiking neural networks. Frontiers in neuroscience 11, 693 (2017)
- 64. Park, S., Kim, S., Na, B., Yoon, S.: T2fsnn: Deep spiking neural networks with time-to-first-spike coding. arXiv preprint arXiv:2003.11741 (2020)
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A.: Automatic differentiation in pytorch. In: NIPS-W (2017)
- Pham, H., Guan, M., Zoph, B., Le, Q., Dean, J.: Efficient neural architecture search via parameters sharing. In: International Conference on Machine Learning. pp. 4095–4104. PMLR (2018)
- Raghu, M., Poole, B., Kleinberg, J., Ganguli, S., Sohl-Dickstein, J.: On the expressive power of deep neural networks. In: international conference on machine learning. pp. 2847–2854. PMLR (2017)
- Rathi, N., Roy, K.: Diet-snn: A low-latency spiking neural network with direct input encoding and leakage and threshold optimization. IEEE Transactions on Neural Networks and Learning Systems (2021)
- Rathi, N., Srinivasan, G., Panda, P., Roy, K.: Enabling deep spiking neural networks with hybrid conversion and spike timing dependent backpropagation. arXiv preprint arXiv:2005.01807 (2020)
- Real, E., Aggarwal, A., Huang, Y., Le, Q.V.: Regularized evolution for image classifier architecture search. In: Proceedings of the aaai conference on artificial intelligence. vol. 33, pp. 4780–4789 (2019)
- Roy, K., Jaiswal, A., Panda, P.: Towards spike-based machine intelligence with neuromorphic computing. Nature 575(7784), 607–617 (2019)
- Rueckauer, B., Lungu, I.A., Hu, Y., Pfeiffer, M., Liu, S.C.: Conversion of continuous-valued deep networks to efficient event-driven networks for image classification. Frontiers in neuroscience 11, 682 (2017)
- Sengupta, A., Ye, Y., Wang, R., Liu, C., Roy, K.: Going deeper in spiking neural networks: Vgg and residual architectures. Frontiers in neuroscience 13, 95 (2019)

- Shrestha, S.B., Orchard, G.: Slayer: Spike layer error reassignment in time. arXiv preprint arXiv:1810.08646 (2018)
- 75. Shu, Y., Wang, W., Cai, S.: Understanding architectures learnt by cell-based neural architecture search. arXiv preprint arXiv:1909.09569 (2019)
- Siems, J., Zimmer, L., Zela, A., Lukasik, J., Keuper, M., Hutter, F.: Nas-bench-301 and the case for surrogate benchmarks for neural architecture search. arXiv preprint arXiv:2008.09777 (2020)
- 77. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. ICLR (2015)
- Tan, M., Chen, B., Pang, R., Vasudevan, V., Sandler, M., Howard, A., Le, Q.V.: Mnasnet: Platform-aware neural architecture search for mobile. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2820–2828 (2019)
- Venkatesha, Y., Kim, Y., Tassiulas, L., Panda, P.: Federated learning with spiking neural networks. arXiv preprint arXiv:2106.06579 (2021)
- Wu, B., Dai, X., Zhang, P., Wang, Y., Sun, F., Wu, Y., Tian, Y., Vajda, P., Jia, Y., Keutzer, K.: Fbnet: Hardware-aware efficient convnet design via differentiable neural architecture search. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10734–10742 (2019)
- 81. Wu, H., Zhang, Y., Weng, W., Zhang, Y., Xiong, Z., Zha, Z.J., Sun, X., Wu, F.: Training spiking neural networks with accumulated spiking flow. ijo 1(1) (2021)
- Wu, J., Chua, Y., Zhang, M., Li, G., Li, H., Tan, K.C.: A tandem learning rule for effective training and rapid inference of deep spiking neural networks. arXiv e-prints pp. arXiv-1907 (2019)
- Wu, J., Xu, C., Zhou, D., Li, H., Tan, K.C.: Progressive tandem learning for pattern recognition with deep spiking neural networks. arXiv preprint arXiv:2007.01204 (2020)
- Wu, Y., Deng, L., Li, G., Zhu, J., Shi, L.: Spatio-temporal backpropagation for training high-performance spiking neural networks. Frontiers in neuroscience 12, 331 (2018)
- Wu, Y., Deng, L., Li, G., Zhu, J., Xie, Y., Shi, L.: Direct training for spiking neural networks: Faster, larger, better. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 33, pp. 1311–1318 (2019)
- Xie, S., Zheng, H., Liu, C., Lin, L.: Snas: stochastic neural architecture search. arXiv preprint arXiv:1812.09926 (2018)
- Xiong, H., Huang, L., Yu, M., Liu, L., Zhu, F., Shao, L.: On the number of linear regions of convolutional neural networks. In: International Conference on Machine Learning. pp. 10514–10523. PMLR (2020)
- Xu, J., Zhao, L., Lin, J., Gao, R., Sun, X., Yang, H.: Knas: green neural architecture search. In: International Conference on Machine Learning. pp. 11613–11625. PMLR (2021)
- Xu, L., Guan, Y., Jin, S., Liu, W., Qian, C., Luo, P., Ouyang, W., Wang, X.: Vipnas: Efficient video pose estimation via neural architecture search. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16072–16081 (2021)
- 90. Yan, Z., Dai, X., Zhang, P., Tian, Y., Wu, B., Feiszli, M.: Fp-nas: Fast probabilistic neural architecture search. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15139–15148 (2021)
- 91. Yang, T.J., Liao, Y.L., Sze, V.: Netadaptv2: Efficient neural architecture search with fast super-network training and architecture optimization. In: Proceedings

20 Y. Kim, Y. Li, H. Park, V. Yeshwanth, P. Panda.

of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2402–2411 (2021)

- 92. Yang, Y., You, S., Li, H., Wang, F., Qian, C., Lin, Z.: Towards improving the consistency, efficiency, and flexibility of differentiable neural architecture search. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6667–6676 (2021)
- 93. Yang, Z., Wang, Y., Chen, X., Guo, J., Zhang, W., Xu, C., Xu, C., Tao, D., Xu, C.: Hournas: Extremely fast neural architecture search through an hourglass lens. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10896–10906 (2021)
- 94. Yao, M., Gao, H., Zhao, G., Wang, D., Lin, Y., Yang, Z., Li, G.: Temporal-wise attention spiking neural networks for event streams classification. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 10221–10230 (2021)
- Ying, C., Klein, A., Christiansen, E., Real, E., Murphy, K., Hutter, F.: Nas-bench-101: Towards reproducible neural architecture search. In: International Conference on Machine Learning. pp. 7105–7114. PMLR (2019)
- Yousefzadeh, A., Stromatias, E., Soto, M., Serrano-Gotarredona, T., Linares-Barranco, B.: On practical issues for stochastic stdp hardware with 1-bit synaptic weights. Frontiers in neuroscience 12, 665 (2018)
- 97. Zeng, D., Huang, Y., Bao, Q., Zhang, J., Su, C., Liu, W.: Neural architecture search for joint human parsing and pose estimation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 11385–11394 (2021)
- Zhang, W., Li, P.: Spike-train level backpropagation for training deep recurrent spiking neural networks. arXiv preprint arXiv:1908.06378 (2019)
- Zhang, W., Li, P.: Temporal spike sequence learning via backpropagation for deep spiking neural networks. arXiv preprint arXiv:2002.10085 (2020)
- 100. Zhang, X., Huang, Z., Wang, N., Xiang, S., Pan, C.: You only search once: Single shot neural architecture search via direct sparse optimization. IEEE Transactions on Pattern Analysis and Machine Intelligence 43(9), 2891–2904 (2020)
- Zhang, X., Xu, H., Mo, H., Tan, J., Yang, C., Wang, L., Ren, W.: Dcnas: Densely connected neural architecture search for semantic image segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13956–13967 (2021)
- 102. Zhang, X., Hou, P., Zhang, X., Sun, J.: Neural architecture search with random labels. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10907–10916 (2021)
- 103. Zhao, Y., Wang, L., Tian, Y., Fonseca, R., Guo, T.: Few-shot neural architecture search. In: International Conference on Machine Learning. pp. 12707–12718. PMLR (2021)
- 104. Zheng, H., Wu, Y., Deng, L., Hu, Y., Li, G.: Going deeper with directly-trained larger spiking neural networks. arXiv preprint arXiv:2011.05280 (2020)
- Zhong, Z., Yan, J., Wu, W., Shao, J., Liu, C.L.: Practical block-wise neural network architecture generation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2423–2432 (2018)
- 106. Zoph, B., Le, Q.V.: Neural architecture search with reinforcement learning. arXiv preprint arXiv:1611.01578 (2016)
- 107. Zoph, B., Vasudevan, V., Shlens, J., Le, Q.V.: Learning transferable architectures for scalable image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8697–8710 (2018)