

LitE-SNN: Designing Lightweight and Efficient Spiking Neural Network through Spatial-Temporal Compressive Network Search and Joint Optimization

Qianhui Liu¹, Jiaqi Yan², Malu Zhang³, Gang Pan^{2*} and Haizhou Li^{4,5,1*}

¹National University of Singapore, Singapore

²Zhejiang University, China

³University of Electronic Science and Technology of China, China

⁴The Chinese University of Hong Kong, Shenzhen (CUHK-Shenzhen), China

⁵Shenzhen Research Institute of Big Data, China

qhliu@nus.edu.sg, jiaqi_yan@zju.edu.cn, maluzhang@uestc.edu.cn, gpan@zju.edu.cn, haizhouli@cuhk.edu.cn

Abstract

Spiking Neural Networks (SNNs) mimic the information-processing mechanisms of the human brain and are highly energy-efficient, making them well-suited for low-power edge devices. However, the pursuit of accuracy in current studies leads to large, long-timestep SNNs, conflicting with the resource constraints of these devices. In order to design lightweight and efficient SNNs, we propose a new approach named LitE-SNN that incorporates both spatial and temporal compression into the automated network design process. Spatially, we present a novel Compressive Convolution block (CompConv) to expand the search space to support pruning and mixed-precision quantization. Temporally, we are the first to propose a compressive timestep search to identify the optimal number of timesteps under specific computation cost constraints. Finally, we formulate a joint optimization to simultaneously learn the architecture parameters and spatial-temporal compression strategies to achieve high performance while minimizing memory and computation costs. Experimental results on CIFAR-10, CIFAR-100, and Google Speech Command datasets demonstrate our proposed LitE-SNNs can achieve competitive or even higher accuracy with remarkably smaller model sizes and fewer computation costs.

1 Introduction

Spiking Neural Networks (SNNs) have gained great attention in recent years [Liu *et al.*, 2020; Zhang *et al.*, 2021; Hu *et al.*, 2021; Hu *et al.*, 2023; Wang *et al.*, 2023; Wei *et al.*, 2024] due to their ability to mimic the information-processing mechanisms of the human brain. They use the timing of the signals (spikes) to communicate between neuronal units. A unit in SNNs is only active when it receives or emits a spike, enabling event-driven processing and high energy efficiency. The received spikes are weight-accumulated

into the membrane potential using only accumulate (AC) operations, which consume much lower energy than the standard energy-intensive multiply-accumulate (MAC) strategy in Artificial Neural Networks (ANNs) [Farabet *et al.*, 2012]. The processing mechanism of SNN can be realized in neuro-morphic hardware, such as Truenorth [Merolla *et al.*, 2014], Loihi [Davies *et al.*, 2018], Darwin [Ma *et al.*, 2024], etc.

The energy efficiency advantages of SNNs make them well-suited for deployment on low-power edge devices. These devices typically have significant resource constraints, such as limited memory capacity and computing resources. Yet, achieving high accuracy with SNN often requires large networks and extended processing timesteps, as they provide superior feature learning and more nuanced data representation. However, this approach increases the demands on memory and computation, which poses challenges for deployment on resource-constrained devices and undermines the potential advantages of SNNs.

Neural architecture search (NAS) aims to automate the neural network design under specific resource constraints. Recent studies have applied NAS to SNN to enable flexible and effective network design and have demonstrated better performance compared to manually designed SNN architectures. However, most of these works do not consider resource constraints in the search process [Na *et al.*, 2022; Che *et al.*, 2022] and only aim at improving the accuracy. While Kim *et al.* consider the number of spikes, thereby reducing communication overhead [Kim *et al.*, 2022a], this metric alone is insufficient for a comprehensive evaluation of resource consumption. To better evaluate resource utilization, we use model size for assessing memory footprint, which is also related to memory access energy, and computation complexity for measuring computational cost. These metrics offer more direct and comprehensive indicators of memory and computing resource demands [Zhang *et al.*, 2016].

To design a lightweight and efficient SNN characterized by small model size and low computation complexity as well as high performance, we incorporate compression within the NAS framework. Existing model compression techniques, including pruning and quantization, reduce the number of parameters and the number of bits used to represent each pa-

*Corresponding author

parameter, respectively, benefiting memory footprint, power efficiency and computational complexity. However, existing SNN pruning and quantization works are all based on fixed architectures, leaving open the question of how to integrate them into evolving SNN architectures in search process. In addition to spatial domain, SNNs uniquely feature a temporal domain that decides time iterations (also known as timesteps). By compressing these iterations, we can directly reduce computational operations and save energy. As current SNN-based NAS methods all originate from ANNs and lack temporal considerations, the potential of implementing the compressive timesteps in NAS has not yet been exploited.

In this paper, we introduce a novel approach named LitE-SNN that incorporates both spatial and temporal compression into the automated network design process, aimed at designing SNNs with compact models and lower computational complexity. Spatially, we propose a Compressive Convolution block (CompConv) that expands the search space to incorporate pruning and quantization. Considering the variable synaptic formats in neuromorphic hardware, CompConv supports mixed-precision quantization to make the network design more flexible and efficient. Meanwhile, we utilize the shared weights and shared pruning masks to reduce the computation introduced by mixed-precision. Temporally, we propose a compressive timestep search that is tailored for the temporal domain of SNNs. Given the computational cost constraint, the model can find the optimal timesteps in accordance with the architecture and spatial compression. Finally, we formulate a multi-objective joint optimization that simultaneously determines the network architecture and compression strategies. This joint optimization ensures a cohesive optimization of these factors, avoiding local optimal and suboptimal global solutions resulting from multiple single-objective optimizations. We conduct the experiments on two image datasets CIFAR-10, CIFAR-100 and a speech dataset Google Speech Command. Experimental results show that our proposed LitE-SNNs achieve competitive accuracy with remarkably smaller model sizes and fewer computation costs. Furthermore, we validate the effectiveness of our LitE-SNNs in balancing accuracy with resource cost and the superiority of our joint optimization. Finally, we conduct energy analysis to further confirm the energy efficiency of LitE-SNNs.

2 Related Work

This section reviews current developments in neural architecture search (NAS) and model compression. NAS aims to automatically design neural architectures that achieve optimal performance using limited resources. Model compression aims to achieve lightweight neural networks with comparable performance, including pruning and quantization, etc.

2.1 Neural Architecture Search

In ANNs, NAS networks have surpassed manually designed architectures on many tasks such as image classification [Liu *et al.*, 2019; Cai *et al.*, 2019]. In the SNN domain, [Na *et al.*, 2022] presented a spike-aware NAS framework involving direct supernet training and an evolutionary search algorithm with spike-aware fitness. However, this approach is

confined to exploring a few predefined blocks, limiting the potential to discover optimal designs beyond them. [Kim *et al.*, 2022a] selected the architecture that can represent diverse spike activation patterns across different data samples. [Che *et al.*, 2022] proposed a spike-based differentiable hierarchical search based on the DARTS [Liu *et al.*, 2019] framework. The SNNs designed by these two works are more flexible and have better accuracy, however, their focus is solely on the accuracy aspect, neglecting critical factors like resource costs. Moreover, they do not consider the time domain during the search process, which is an important aspect of SNN.

2.2 Model Compression

Pruning involves removing unnecessary connections to reduce the number of parameters and computational needs. In ANNs, many existing studies employed important scores-based pruning [Sehwag *et al.*, 2020], which involves training importance scores for each weight and masking less important weights. In SNNs, [Chen *et al.*, 2021] proposed a gradient-based rewiring method to learn the connectivity and weight parameters of SNNs. [Chen *et al.*, 2022] addressed the pruning of deep SNNs by modeling the state transition of dendritic spine. [Kim *et al.*, 2022b] introduced the lottery ticket hypothesis-based pruning for sparse deep SNNs.

Quantization reduces the number of bits used to represent each parameter to reduce the memory footprint and computational complexity. Some ANN-based works [Cai *et al.*, 2017; Cai and Vasconcelos, 2020] set the quantizer based on the weight and activation distribution. [Rueckauer *et al.*, 2017] converted the binary ANN to obtain a binary SNN. [Srinivasan and Roy, 2019] used STDP to learn the one-bit weight synapses. However, these SNN methods apply uniform bit-width, neglecting the different sensitivities of different filters. [Lui and Neftci, 2021] employed a layer-wise Hessian trace analysis for mixed-precision quantization, allocating the bit-width per layer, but limited to a three-layer shallow network.

There are several works that consider both pruning and quantization. [Deng *et al.*, 2021] adapted alternating direction method of multipliers (ADMM) optimization with spatio-temporal backpropagation for SNN pruning and quantization. However, this approach is performed sequentially and may lead to suboptimal results, e.g., the best network architecture for the dense and full-precision model is not necessarily the optimal one after pruning and quantization [Wang *et al.*, 2020]. [Rathi *et al.*, 2018] presented a joint pruned and quantized SNN with STDP-based learning, but it is limited to a fixed and shallow (two layers) structure. Current SNN research lacks a comprehensive approach that simultaneously considers the neural architecture design, pruning and quantization. Therefore, we need a solution to jointly optimize these aspects. Similar work in ANN needs to train a large supernet and distill a number of smaller sub-networks, followed by training the accuracy predictors and evolutionary search [Wang *et al.*, 2020]. The whole training process takes 100 GPU days on V100 GPU. Considering that SNNs usually require more training time than ANNs [Kim *et al.*, 2022a], it is not practical to utilize this approach for SNNs.

In the time domain, [Li *et al.*, 2023] dynamically determines the number of timesteps during inference on an

input-dependent basis. While this work effectively reduces timesteps, it is limited to the inference stage and cannot be optimized together with other training parameters, which also makes it unsuitable for use in NAS. In this paper, we will jointly optimize the pruning, mixed-precision quantization, and timestep compression through the NAS framework.

3 Method

In this section, we first introduce the preliminary knowledge of spiking neurons and architecture search space, then present our proposed spatial and temporal compression search, and finally explain the spatial-temporal joint optimization.

3.1 Preliminary

We choose a widely used spiking neural model Leaky Integrate-and-Fire (LIF) [Zheng *et al.*, 2021] to describe the neural behavior, which can be written as:

$$u_i^{t,n} = \tau_{decay} u_i^{t-1,n} (1 - y_i^{t-1,n}) + \sum_j W_{i,j}^n y_j^{t,n-1} \quad (1)$$

where τ_{decay} denotes the membrane decay constant, W_{ij} is the synaptic weight between the j -th neuron (or feature map) and the i -th neuron (or feature map). If the membrane potential at timestamp t of the neuron in n -th layer $u_i^{t,n}$ is larger than the threshold V_{th} , the output spike $y_i^{t,n}$ is set to 1; otherwise it is set to 0. Our method builds upon the search space of spikeDHS [Che *et al.*, 2022], which is a spiking-based differentiable hierarchical search framework that can find best-performed architectures for SNNs. The search space consists of a number of cells and each cell (depicted in Figure 1(a)) is defined as a repeated and searchable unit with N nodes (depicted in Figure 1(b)), $\{x_i\}_N$. Each cell receives input from two previous cells and forms its output by concatenating all outputs of its nodes. Each node can be described by:

$$x_j = f\left(\sum_{i < j} o^{i,j}(x_i)\right) \quad (2)$$

where f is a spiking neuron taking the sum of all operations as input, $o^{(i,j)}$ is the operation associated with the directed edge connecting node i and j . During search, each edge is represented by a weighted average of candidate operations. The information flow connecting node i and node j becomes:

$$\bar{o}^{(i,j)}(x) = \sum_{o \in \mathcal{O}^{i,j}} \frac{\exp(\alpha_o^{(i,j)})}{\sum_{o \in \mathcal{O}^{i,j}} \exp(\alpha_o^{(i,j)})} \cdot o(x) \quad (3)$$




Figure 1: Hierarchical search space. (a) Cell structure search space with 3 nodes. (b) Within one node, two candidate operations are summed filtered by a sign function.

where $\mathcal{O}^{i,j}$ denotes the operation space on edge (i, j) and $\alpha_o^{(i,j)}$ is the weight of operation o , which is a trainable continuous variable. After search, a discrete architecture is selected by replacing each mixed operation $\bar{o}^{(i,j)}$ with the most likely operation $o^{i,j}$ that has $\max_{o \in \mathcal{O}^{(i,j)}} \alpha_o^{(i,j)}$.

3.2 Compressive Convolution (CompConv)

To design lightweight and efficient SNNs, in this work, we incorporate compression techniques into the neural architecture search to take advantage of the efficiency and flexibility of automatic network design. To this end, we propose a novel compressive convolutional block (CompConv) to replace the standard convolution (green block in Figure 1) to enable the search of compressive SNNs. Neuromorphic hardware like Loihi is capable of supporting sparse kernels and variable synaptic formats [Davies *et al.*, 2018; Davies and others, 2021], which inspired us to utilize weight pruning and mixed-precision quantization for compression.

As shown in Figure 2(a), a **naive solution** to integrate mixed-precision quantization with pruning is to organize each bit-width quantization operation and the corresponding pruning mask into a separate branch and update each branch individually, which can be described

$$W_q^{(j)} = Q_j(W^{(j)}) \quad (4)$$

$$e^{(j)}(W) = W_q^{(j)} \odot \text{mask}^{(j)} \quad (5)$$

$$\bar{e}(W) = \sum_j^{\text{len}(B)} \frac{\exp(\beta^{(j)})}{\sum_j^{\text{len}(B)} \exp(\beta^{(j)})} e^{(j)}(W) \quad (6)$$

$$W_{\text{output}} = \bar{e}(W) \quad (7)$$

where $W^{(j)}$ represents the weight tensor of j -th branch, Q_j is a quantization function that follows [Cai *et al.*, 2017], which quantizes the input value by the j -th candidate bitwidth. $\text{mask}^{(j)}$ is a binary matrix where each element indicates whether the weights have been pruned (assigned a 0) or retained (assigned a 1) of j -th branch. The pruning method follows the [Sehwag *et al.*, 2020]. $W_q^{(j)}$ and $e^{(j)}(W)$ represent the quantized weights and operation of j -th branch respectively. $\beta^{(j)}$ is the weight of j -th operation and B denotes the bit-width candidate space. However, this approach significantly increases the computation, as we need to update the weights and mask for each individual branch. Additionally, it has been found in [Cai and Vasconcelos, 2020] that branches with low β values receive few gradients, potentially leading to the under-trained of weights and mask.

In contrast, our CompConv provides a **more efficient solution** by sharing the weights and pruning mask among all branches within a block. As depicted in Figure 2(b), we first quantize W with different bit-width in different branches to get W_q and then perform pruning mask on the weighted average of W_q . The Equation (4-7) can be rewritten as

$$W_q^{(j)} = Q_j(W) \quad (8)$$

$$e^{(j)}(W) = W_q^{(j)} \quad (9)$$

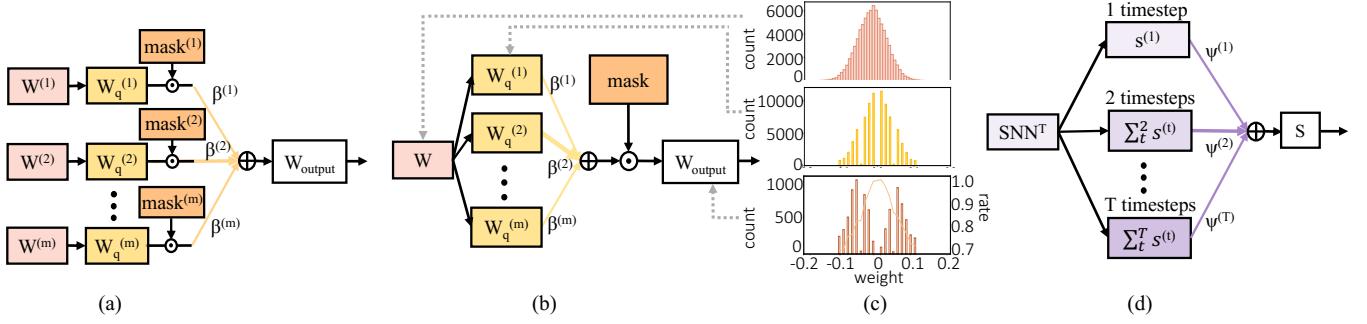


Figure 2: (a) A naive solution that integrates mixed-quantization with pruning. (b) Our proposed CompConv for a more efficient integration. (c) Visualization of the distribution of W , $\bar{e}(W)$, and W_{output} with pruning rates. (d) Our proposed timestep search solution.

$$\bar{e}(W) = \sum_j^{\text{len}(B)} \frac{\exp(\beta^{(j)})}{\sum_j^{\text{len}(B)} \exp(\beta^{(j)})} e^{(j)}(W) \quad (10)$$

$$W_{output} = \bar{e}(W) \odot mask \quad (11)$$

where we remove the $W^{(j)}$ and $mask^{(j)}$ in each branch and utilize universal shared W and $mask$ instead. This modification allows the gradient to be fully applied to the W and $mask$, eliminating the possibility of underfitting and simplifying the search space. Overall, our proposed CompConv offers a more efficient solution for achieving compression with mixed-precision quantization and pruning in the search space of automatic lightweight SNN designs.

Figure 2(c) illustrates the distributions of W , $\bar{e}(W)$, and W_{output} with pruning. W displays a symmetric, continuous distribution with a diminishing central peak. $\bar{e}(W)$, quantized at 2-bit and 4-bit precision, follows a similar discrete pattern. The pruning rate curve indicates that weights with smaller absolute values are pruned more for their lower importance, resulting in W_{output} having higher values on the sides and lower at the center and far sides.

3.3 Compressive Timestep Search

The selection of timesteps in SNNs is crucial, as it influences the network’s ability to process temporal information. A longer timestep can improve accuracy but comes with increased computational cost due to the necessity for more time iterations. Conversely, a shorter timestep may reduce computational cost but potentially at the expense of temporal detail and accuracy. Recognizing this, our work focuses on incorporating the timestep selection into the search process to automatically balance computational cost with performance.

To assess the computational operations associated with different timesteps, we follow the convention of the neuromorphic computing community to use the total synaptic operations (SynOps) metric [Wu *et al.*, 2021], defined as:

$$SynOps = (1 - p\%)k_h k_w C_{in} C_{out} H W S \quad (12)$$

where p is the pruning rate, k_h and k_w are the kernel height and width, C_{in} and C_{out} the input and output channel sizes, H and W the height and width of the feature map. S is the average spike rate, which is originally defined as

$$S = \sum_t^T s^{(t)} \quad (13)$$

where T is the total time length, $s^{(t)}$ indicates the average spike rate per neuron at the t -th timestep. We observe that the number of timesteps affects the spike rate S and, consequently, the network’s computational cost. As shown in Figure 2(d), to enable this search, we consider a range of candidate timestep counts $\{1, 2, \dots, T\}$ and modify the Equation (13) by introducing trainable weights for each timestep count:

$$S = \sum_t^T \left(\frac{\exp(\psi^{(t)})}{\sum_t^T \exp(\psi^{(t)})} \sum_{t'}^t s^{(t')} \right) \quad (14)$$

where $\psi^{(t)}$ is the weight of having t timesteps, $\frac{\exp(\psi^{(t)})}{\sum_t^T \exp(\psi^{(t)})}$ is the relative weight with the total sum equaling 1. For illustration, consider a scenario where the total time length T is 6 with relative weights for each timestep $t \in \{1, 2, 3, 4, 5, 6\}$ as $[0, 0, 0, 0, 0.6, 0.4]$. Here, the spike rate S includes the sum of spike rates for the first 5 timesteps plus 40% from the 6th. If weights change to $[0, 0, 0, 0, 0, 1]$, S equals the sum across all timesteps, matching Equation (13). The trained weights ψ evaluate the impact of timestep selection on network performance. We will introduce how to balance network performance against computational cost in Section 3.4

3.4 Loss Function and Joint Optimization Algorithm

The objective of LitE-SNN is to identify SNNs that are both lightweight and efficient, achieving high performance while minimizing memory and computation costs. This can naturally be formulated as a multi-objective optimization problem. We formulate the following loss function:

$$Loss = Loss_{CE} + \lambda_1 Loss_{MEM} + \lambda_2 Loss_{COMP} \quad (15)$$

where the coefficients λ_1 and λ_2 serve as regularization parameters that control the trade-off among $Loss_{CE}$, $Loss_{MEM}$ and $Loss_{COMP}$. The first term $Loss_{CE}$ represents the cross-entropy (CE) loss, responsible for evaluating the predictive accuracy of the SNN. Following [Che *et al.*, 2022], the original definition of $Loss_{CE}$ is:

$$Loss_{CE} = CE^{(T)} = CE \left(\frac{1}{T} \sum_{t=1}^T O(t), y \right), \quad (16)$$

where $O(t)$ denotes the output of the SNN and y represents the target label. But as $Loss_{CE}$ is also influenced by the number of timesteps, we rewrite it as:

$$Loss_{CE} = \sum_t^T \frac{\exp(\psi^{(t)})}{\sum_t^T \exp(\psi^{(t)})} \cdot CE^{(t)} \quad (17)$$

The second term $Loss_{MEM}$ quantifies the memory usage of the network, which is measured by the model size:

$$Loss_{MEM} = b_w k_h k_w C_{in} C_{out} (1 - p\%) \quad (18)$$

where b_w is the bit-width of the network weights. During search, the b_w is a weighted average of candidate bit-width and can be calculated as

$$b_w = \sum_j^{\text{len}(B)} \frac{\exp(\beta^{(j)})}{\sum_j^{\text{len}(B)} \exp(\beta^{(j)})} \cdot q^{(j)} \quad (19)$$

where q^j denotes the j -th bitwidth candidate. The third term $Loss_{COMP}$ assesses the computational cost, which is affected not only by the SynOps but also by the operational precision [Cai and Vasconcelos, 2020]. Therefore, we define this term as bit-SynOps, reflecting the computational complexity of the model. The corresponding equation is:

$$Loss_{COMP} = \text{Bit-SynOps} = b_w \text{SynOps} \quad (20)$$

The joint optimization comprises two steps for each iteration. First, we conduct the first forward pass and update the weights, which constitute the foundational components of the network. Then, a second forward pass is conducted using the updated weights and we simultaneously update β and $mask$ within our CompConv, as well as the ψ and architecture parameters α . Both two steps apply the loss function in Equation (15). Our joint optimization ensures cohesive and efficient optimization, avoiding local optimal and suboptimal global solutions resulting from multiple single-objective optimizations [Wang *et al.*, 2020].

After search, we decode the cell structure (in Figure 1(a)) by retaining the two strongest incoming edges for each node and select the node operation (in Figure 1(b)) with the strongest edge, where the strength of edge is determined by α . The bit-width (in Figure 2(b)) chosen for each cell is decoded by the branch with the highest weight β . We also select the number of timesteps with the highest weight ψ (in Figure 2(d)). During the retraining, we use the auxiliary loss as in [Che *et al.*, 2022]. The weights and masks are retrained to adapt to single-branch network. We adopt the surrogate gradient to direct train the SNN and use the Dspike function [Li *et al.*, 2021] to approximate the derivative of spike activity.

4 Experiments

This section compares our model’s performance and efficiency with current benchmarks and optimization methods and evaluates the effectiveness of parameters λ_1 and λ_2 .

4.1 Experimental Settings

Datasets. We evaluate our proposed LitE-SNN on two image datasets: CIFAR-10 [Krizhevsky *et al.*, 2009], CIFAR-100

Layer	CIFAR	GSC
	channel \times feature size	channel \times feature size
stem	$(C_{\text{CIFAR}}) \times (32 \times 32)$	$(C_{\text{GSC}}) \times (40 \times 98)$
Cell1	$(C_{\text{CIFAR}}) \times (32 \times 32)$	$(C_{\text{GSC}}) \times (40 \times 98)$
Cell2	$(C_{\text{CIFAR}}) \times (32 \times 32)$	$(C_{\text{GSC}}) \times (40 \times 98)$
Cell3	$(2 \times C_{\text{CIFAR}}) \times (16 \times 16)$	$(2 \times C_{\text{GSC}}) \times (20 \times 49)$
Cell4	$(2 \times C_{\text{CIFAR}}) \times (16 \times 16)$	$(2 \times C_{\text{GSC}}) \times (20 \times 49)$
Cell5	$(2 \times C_{\text{CIFAR}}) \times (16 \times 16)$	$(4 \times C_{\text{GSC}}) \times (10 \times 25)$
Cell6	$(4 \times C_{\text{CIFAR}}) \times (8 \times 8)$	$(4 \times C_{\text{GSC}}) \times (10 \times 25)$
Cell7	$(4 \times C_{\text{CIFAR}}) \times (8 \times 8)$	-
Cell8	$(4 \times C_{\text{CIFAR}}) \times (8 \times 8)$	-
Pooling	$(4 \times C_{\text{CIFAR}}) \times (1 \times 1)$	$(4 \times C_{\text{GSC}}) \times (1 \times 1)$
FC	10 / 100	12

Table 1: Network backbone for CIFAR and GSC datasets.

[Krizhevsky *et al.*, 2009]. We also choose an audio dataset Google Speech Command (GSC) [Warden, 2018] which represents the keyword-spotting task commonly deployed on edge devices for quick responses.

CIFAR-10 has 60,000 images of 10 classes with a size of 32×32 . Among them, there are 50,000 training images and 10,000 testing images. CIFAR-100 has the same configurations as CIFAR-10, except it contains 100 classes. We split the dataset into 9,000 training samples and 1,000 test samples. The pre-processing method we use on these two datasets is the same as [Che *et al.*, 2022].

GSC has 35 words spoken by 2,618 speakers. Following the pre-processing method used in [Yang *et al.*, 2022], we split the dataset into 12 classes, which include 10 keywords and two additional classes. We apply Mel Frequency Cepstrum Coefficient (MFCC) to extract acoustic features. The sampling frequency is 16kHz, the frame length and shift are set to 30ms and 10ms, and the filter channel is defined as 40.

Hyperparameters. We implement our LitE-SNN using Pytorch on NVIDIA GeForce RTX 3080 (10G) GPUs. During the architecture search, we conduct 50 epochs with a batch size of 30. We use the SGD optimizer with momentum 0.9 and a learning rate of 0.025 to update network weights w and employ the Adam optimizer with a learning rate of $3e^{-4}$ to update β , $mask$, and α . During the retraining, we train for 200 epochs with a batch size of 50, using the SGD optimizer with momentum 0.9 and a cosine learning rate of 0.025. The neuron parameters τ_{decay} and V_{reset} are set to 0.2 and 0 respectively. The threshold V_{th} is initialized to 0.5. The temperature parameter in Dspike function is set to 3 and the settings of other parameters follow [Che *et al.*, 2022].

Architecture. The architecture backbone is reported in Table 1. The stem layer has a structure of Conv-Spike-BN with a kernel size of 3×3 . The initial channel counts are denoted as C_{CIFAR} and C_{GSC} for CIFAR and GSC datasets, with respective values of 48 and 16 during the search. For the rest of the network, we employ 8 searchable cells on CIFAR dataset and 6 searchable cells on GSC dataset. The final two layers are a global pooling and a fully-connected layer. Each of the searchable cells has 4 nodes with each node containing two candidate operations: Conv 3×3 and skip. In this paper, the bit-width candidate space is set as $\{1, 2, 4\}$. For image datasets, we search on CIFAR-10 and then retrain on target datasets including CIFAR-10 and CIFAR-100. For audio dataset GSC, we use the architecture searched on its

Model	CIFAR-10					CIFAR-100				
	Acc. (%)	Model Size (MB)	Bit-SynOps (M)	Bitwidth. (b)	#Timesteps	Acc. (%)	Model Size (MB)	Bit-SynOps (M)	Bitwidth. (b)	#Timesteps
[Rueckauer <i>et al.</i> , 2017]	83.35	0.27	-	1	4	-	-	-	-	-
[Srinivasan and Roy, 2019]	66.23	2.10	-	1	25	-	-	-	-	-
[Chen <i>et al.</i> , 2021]	92.54 92.50	41.73 17.68	847 638	32 (FP) 32 (FP)	8 8	69.36 67.47	11.60 2.56	1179 882	32 (FP) 32 (FP)	8 8
[Deng <i>et al.</i> , 2021]	87.84 87.59	8.62 2.87	- -	3 3	8 8	57.83 55.95	5.75 2.87	- -	3 1	8 8
[Kim <i>et al.</i> , 2022b]	93.50 93.46	12.02 2.88	950 560	32 (FP) 32 (FP)	5 5	71.45 71.00	12.02 2.88	663 517	32 (FP) 32 (FP)	5 5
[Chen <i>et al.</i> , 2022]	92.49 90.21	3.28 1.10	1290 1245	32 (FP) 32 (FP)	8 8	-	-	-	-	-
[Kim <i>et al.</i> , 2022a]	94.12 93.73	199.72 168.24	16504 13203	32 (FP) 32 (FP)	8 5	73.04 70.06	82.84 77.60	39243 27735	32 (FP) 32 (FP)	5 5
[Na <i>et al.</i> , 2022]	93.15 92.54	83.68 21.76	18739 12787	32 (FP) 32 (FP)	8 8	69.16	42.76	19547	32 (FP)	8
[Che <i>et al.</i> , 2022]	95.50 95.36	56.00 56.00	41781 40564	32 (FP) 32 (FP)	6 6	76.25 76.03	56.00 48.00	51577 50270	32 (FP) 32 (FP)	6 6
[Xu <i>et al.</i> , 2023]	93.41	44.68	-	32 (FP)	4	-	-	-	-	-
[Shen <i>et al.</i> , 2023]	91.09	25.12	-	32 (FP)	2	73.48	25.12	-	32 (FP)	4
[Yan <i>et al.</i> , 2024]	94.64 94.27	93.88 64.72	13184 8640	32 (FP) 32 (FP)	3 3	74.78 73.21	110.20 69.28	16544 11136	32 (FP) 32 (FP)	3 3
Our Work (large)	95.60±0.24	3.60	2863	searched	searched (6)	77.10±0.04	3.62	3590	searched	searched (6)
Our Work (medium)	94.52±0.05	1.23	913	searched	searched (5)	74.51±0.35	1.27	1110	searched	searched (5)
Our Work (small)	91.98±0.15	0.55	298	searched	searched (4)	69.55±0.11	0.59	372	searched	searched (4)

Table 2: Performance comparison with the state-of-the-art lightweight SNNs on CIFAR10 and CIFAR100 datasets.

own. During retraining, C_{CIFAR} is increased to 144 to extract more features from raw data. To compare with state-of-the-art SNNs which span a wide range of model size and Bit-SynOps on CIFAR, we set three scales of LitE-SNNs. Large scale uses the λ_1 , λ_2 and p of 0, 0 and 50; medium scale employs $5e^{-10}$, $5e^{-14}$ and 80; and small scale adopt $1e^{-9}$, $1e^{-13}$ and 90. On GSC, we use λ_1 , λ_2 and p of $1e^{-9}$, $1e^{-13}$ and 70.

4.2 Comparison with Existing SNNs

Table 2¹ presents a comprehensive comparison between our proposed LitE-SNNs and other state-of-the-art lightweight SNNs on CIFAR10 and CIFAR-100 datasets. Our large-scale LitE-SNN achieves state-of-the-art accuracy on CIFAR-10 and CIFAR-100 datasets while maintaining a remarkably compact model size of 3.60MB and 3.62MB and relatively fewer BitSynOps of 2863M and 3590M. We can see that, this compressed SNN achieves superior accuracy compared with the baseline model without compression [Che *et al.*, 2022]. This result reveals the potential of compression to remove noisy and irrelevant connections to enhance the ability to capture valuable information. Our small-scale LitE-SNN has a model size of only 0.55MB and Bit-SynOps of 298M, smaller than all other models except the model size of [Rueckauer *et al.*, 2017] on CIFAR-10 dataset. Nevertheless, it achieves a competitive accuracy of 91.98%, outperforming [Rueckauer *et al.*, 2017] by a big margin of 8.63%. Notice that [Chen *et al.*, 2021], [Kim *et al.*, 2022b] and [Chen *et al.*,


¹The results of comparison SNNs are from original papers, or (if not provided) from our experiments using publicly available code from the paper. If the original paper reported multiple results, we list two results with the highest accuracy in the table.

Model	Acc. (%)	Model Size (MB)	Bit-SynOps (M)	Bitwidth. (b)	#Timesteps
[Yilmaz <i>et al.</i> , 2020]	75.20	0.47	-	32 (FP)	10
[Orchard <i>et al.</i> , 2021]	88.97	0.95	-	32 (FP)	-
[Yang <i>et al.</i> , 2022]	92.90	0.36	768	32 (FP)	32
Our Work	94.14±0.06	0.30	37	searched	searched (3)

Table 3: Performance comparison with other SNNs on Google Speech Command dataset.

2022] also maintain a good performance with relatively small model size and Bit-SynOps. This is because they employ effective pruning methods to reduce the number of parameters. Despite this, under similar model sizes or Bit-SynOps, our LitE-SNNs still achieve higher accuracy. For example, on CIFAR-10 dataset, our medium-scale LitE-SNN achieves 94.52% accuracy with a 1.23MB model, surpassing [Kim *et al.*, 2022b] by 1% while being 1.6MB smaller. The results of our timestep search are also shown in Table 2. From large-scale to small-scale LitE-SNN, we observe that the optimal number of timesteps decreases from 6 to 4. This is because the increased λ_1 and λ_2 values in the search impose stricter resource limits. Consequently, the network tends to select a smaller number of timesteps to minimize calculations.

On the GSC dataset, as reported in Table 3, our model gets an accuracy of 94.14% with a compact model size of 0.3M and Bit-SynOps of only 37M, which outperforms all the comparison methods in terms of accuracy and resource cost. It is worth noting that the model sizes of other compared methods are all less than 1MB. This is because GSC belongs to the keyword-spotting task that tends to run on the portable device. For such tasks, it is important to employ smaller models


 Figure 3: Ablation study of λ_1 and λ_2 .

Model	Acc. (%)	Model Size (MB)	Bit-SynOps (M)	Design Cost (GPU days)
Sequential Optimization				
A + P + Q + T (large)	94.73	3.43	2426	7.04
A + P + Q + T (medium)	92.76	1.40	1001	7.04
A + P + Q + T (small)	85.99	0.73	505	7.04
Joint Optimization				
LitE-SNN (large)	95.60	3.60	2863	5.01
LitE-SNN (medium)	94.52	1.23	913	4.99
LitE-SNN (small)	91.98	0.55	298	4.48

Table 4: Performance comparison of our joint optimization with sequential optimization.

to minimize power consumption and speed up the response.

4.3 Effectiveness of λ_1 and λ_2

We vary λ_1 and λ_2 in Equation (15) and validate its effect on the trade-off between accuracy and resource costs in the LitE-SNN on CIFAR-10 dataset. The other parameters maintain the same setting as our small-scale LitE-SNN.

Figure 3 illustrates that increasing λ_1 and λ_2 leads to stricter resource limits, resulting in a trend of smaller network sizes and fewer Bit-SynOps, but at the expense of accuracy. Notably, λ_1 (the memory coefficient) and λ_2 (the computational cost coefficient) have interconnected effects on model size and Bit-SynOps. For instance, a rise in λ_2 alone, altering the network’s search setting from (1e-9, 0) to (1e-9, 1e-13), impacts both Bit-SynOps and model size. This is because the constraint on Bit-SynOps limits model size, reducing both metrics. Similarly, adjusting the settings from (1e-9, 1e-13) to (5e-9, 1e-13) also decreases both model size and Bit-SynOps. The search results demonstrate that λ_1 and λ_2 can trade-off between accuracy and computational cost, while also coordinating changes in model size and Bit-SynOps.

4.4 Comparison with Sequential Optimization

In this section, we conduct a comparative analysis between our LitE-SNNs and the SNNs optimized using the sequential scheme on CIFAR-10 dataset to validate the effectiveness of our joint optimization scheme. The comparison models are derived by sequentially performing neural architecture search (A), pruning (P), quantization (Q) and timestep selection (T). The methods used in each step follow the corresponding scales of LitE-SNNs.

As shown in Table 4, the three versions of LitE-SNN consistently outperform models that are derived from sequential optimization, achieving higher accuracy with comparable model sizes and Bit-SynOps. Since each step in sequential models applies the same methods as LitE-SNNs, their inferior accuracy and longer design time demonstrate the effectiveness of our joint optimization. It is worth noting that the small-scale sequential model suffers a substantial accuracy

Model	#Addition(M)	#Multiplication(M)	Energy (mJ)
ANN (large)	970.12	970.12	4.46
ANN (medium)	413.53	413.53	1.90
ANN (small)	221.63	221.63	1.02
AutoSNN [Na <i>et al.</i> , 2022]	585.60	28.31	0.63
SpikeDHS [Che <i>et al.</i> , 2022]	1305.65	23.89	1.26
LitE-SNN (large)	705.01	11.94	0.68
LitE-SNN (medium)	271.41	4.78	0.26
LitE-SNN (small)	129.74	2.39	0.13

Table 5: Energy cost comparison for a single forward.

loss. This can be attributed to its high sparsity of 90%, making it sensitive to parameter adjustments. Small changes in parameter values due to subsequent quantization can significantly impact the network. Whereas our joint optimization considers the intertwined effects of various compressive factors, thereby effectively mitigating accuracy degradation.

4.5 Energy Consumption

In this section, we estimate the theoretical energy consumption of our LitE-SNN using the common approach in the neuromorphic community [Rathi and Roy, 2021; Li *et al.*, 2021]. The energy is measured in 45nm CMOS technology [Horowitz, 2014], where the addition operation and multiplication operation cost 0.9 pJ and 3.7 pJ energy respectively. These energy values are based on 32-bit floating-point operations. Since our model uses quantization to a smaller bit width, the reported energy values for our model are overestimates compared to those of other models.

Table 5 presents the energy consumption of our LitE-SNNs and corresponding ANN counterparts with the same network architecture. Our LitE-SNNs with three scales respectively consume over $6\times$ lower energy compared with ANN. Compared with our baseline model SpikeDHS, LitE-SNNs achieve an energy reduction of more than 0.58mJ. We also compare our results with AutoSNN, which focuses on designing energy-efficient SNNs. Its energy consumption is comparable with our large-scale LitE-SNN, larger than our small and medium-scale LitE-SNNs. We acknowledge that energy consumption is not only from computation but also from memory access, a topic beyond the scope of our study as it is more about hardware design. Nevertheless, it’s worth noting that our model’s smaller size and lower bit-width can contribute to reducing access-related energy consumption.

5 Conclusion

In this paper, we propose the LitE-SNN to integrate spatio-temporal compression into the automatic network design. We present the CompConv that expands the search space to support pruning and mixed-precision quantization. Meanwhile, we propose the compressive timestep search to identify the optimal number of timesteps to process SNNs. We also formulate a joint optimization algorithm to learn these compression strategies and architecture parameters. Experimental results show that our LitE-SNN exhibits competitive accuracy, compact model size, and fewer computation costs, making it an attractive choice for deployment on resource-constrained edge devices and opening new possibilities for the practical implementation of SNNs in real-world applications.

Acknowledgements

This work is supported by the Natural Science Foundation of China (No. 61925603, U1909202) and the Key Research and Development Program of Zhejiang Province in China (2020C03004). This work is also supported by IAF, A*STAR, SOITEC, NXP, National University of Singapore under FD-fAbrICS: Joint Lab for FD-SOI Always-on Intelligent & Connected Systems (Award I2001E0053) and the Internal Project Fund from Shenzhen Research Institute of Big Data under Grant T00120220002.

References

- [Cai and Vasconcelos, 2020] Zhaowei Cai and Nuno Vasconcelos. Rethinking differentiable search for mixed-precision neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2349–2358, 2020.
- [Cai et al., 2017] Zhaowei Cai, Xiaodong He, Jian Sun, and Nuno Vasconcelos. Deep learning with low precision by half-wave gaussian quantization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5918–5926, 2017.
- [Cai et al., 2019] Han Cai, Ligeng Zhu, and Song Han. Proxylessnas: Direct neural architecture search on target task and hardware. In *7th International Conference on Learning Representations*, 2019.
- [Che et al., 2022] Kaiwei Che, Luziwei Leng, Kaixuan Zhang, Jianguo Zhang, Qinghu Meng, Jie Cheng, Qinghai Guo, and Jianxing Liao. Differentiable hierarchical and surrogate gradient search for spiking neural networks. *Advances in Neural Information Processing Systems*, 35:24975–24990, 2022.
- [Chen et al., 2021] Yanqi Chen, Zhaofei Yu, Wei Fang, Tiejun Huang, and Yonghong Tian. Pruning of deep spiking neural networks through gradient rewiring. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*, pages 1713–1721, 2021.
- [Chen et al., 2022] Yanqi Chen, Zhaofei Yu, Wei Fang, Zhengyu Ma, Tiejun Huang, and Yonghong Tian. State transition of dendritic spines improves learning of sparse spiking neural networks. In *International Conference on Machine Learning*, pages 3701–3715, 2022.
- [Davies and others, 2021] Mike Davies et al. Taking neuromorphic computing to the next level with loihi2. *Intel Labs’ Loihi*, 2:1–7, 2021.
- [Davies et al., 2018] Mike Davies, Narayan Srinivasa, Tsung-Han Lin, Gautham Chinya, Yongqiang Cao, Sri Harsha Choday, Georgios Dimou, Prasad Joshi, Nabil Imam, Shweta Jain, et al. Loihi: A neuromorphic manycore processor with on-chip learning. *IEEE Micro*, 38(1):82–99, 2018.
- [Deng et al., 2021] Lei Deng, Yujie Wu, Yifan Hu, Ling Liang, Guoqi Li, Xing Hu, Yufei Ding, Peng Li, and Yuan Xie. Comprehensive snn compression using admm optimization and activity regularization. *IEEE Transactions on Neural Networks and Learning Systems*, 2021.
- [Farabet et al., 2012] Clément Farabet, Rafael Paz, Jose Pérez-Carrasco, Carlos Zamarreño-Ramos, Alejandro Linares-Barranco, Yann LeCun, Eugenio Culurciello, Teresa Serrano-Gotarredona, and Bernabe Linares-Barranco. Comparison between frame-constrained fix-pixel-value and frame-free spiking-dynamic-pixel convnets for visual processing. *Frontiers in Neuroscience*, 6:32, 2012.
- [Horowitz, 2014] Mark Horowitz. 1.1 computing’s energy problem (and what we can do about it). In *2014 IEEE international solid-state circuits conference digest of technical papers (ISSCC)*, pages 10–14. IEEE, 2014.
- [Hu et al., 2021] Yangfan Hu, Huajin Tang, and Gang Pan. Spiking deep residual networks. *IEEE Transactions on Neural Networks and Learning Systems*, 34(8):5200–5205, 2021.
- [Hu et al., 2023] Yangfan Hu, Qian Zheng, Xudong Jiang, and Gang Pan. Fast-snn: Fast spiking neural network by converting quantized ann. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.
- [Kim et al., 2022a] Youngeun Kim, Yuhang Li, Hyoungseob Park, Yeshwanth Venkatesha, and Priyadarshini Panda. Neural architecture search for spiking neural networks. In *European Conference on Computer Vision*, pages 36–56. Springer, 2022.
- [Kim et al., 2022b] Youngeun Kim, Yuhang Li, Hyoungseob Park, Yeshwanth Venkatesha, Ruokai Yin, and Priyadarshini Panda. Exploring lottery ticket hypothesis in spiking neural networks. In *European Conference on Computer Vision*, pages 102–120. Springer, 2022.
- [Krizhevsky et al., 2009] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [Li et al., 2021] Yuhang Li, Yufei Guo, Shanghang Zhang, Shikuang Deng, Yongqing Hai, and Shi Gu. Differentiable spike: Rethinking gradient-descent for training spiking neural networks. In *Advances in Neural Information Processing Systems*, pages 23426–23439, 2021.
- [Li et al., 2023] Yuhang Li, Abhishek Moitra, Tamar Geller, and Priyadarshini Panda. Input-aware dynamic timestep spiking neural networks for efficient in-memory computing. In *2023 60th ACM/IEEE Design Automation Conference*, pages 1–6. IEEE, 2023.
- [Liu et al., 2019] Hanxiao Liu, Karen Simonyan, and Yiming Yang. DARTS: differentiable architecture search. In *International Conference on Learning Representations*, 2019.
- [Liu et al., 2020] Qianhui Liu, Gang Pan, Haibo Ruan, Dong Xing, Qi Xu, and Huajin Tang. Unsupervised aer object recognition based on multiscale spatio-temporal features and spiking neurons. *IEEE Transactions on Neural Networks and Learning Systems*, 31(12):5300–5311, 2020.
- [Lui and Neftci, 2021] Hin Wai Lui and Emre Neftci. Hes-sian aware quantization of spiking neural networks. In *International Conference on Neuromorphic Systems 2021*, pages 1–5, 2021.

- [Ma *et al.*, 2024] De Ma, Xiaofei Jin, Shichun Sun, Yitao Li, Xundong Wu, Youneng Hu, Fangchao Yang, Huajin Tang, Xiaolei Zhu, Peng Lin, et al. Darwin3: A large-scale neuromorphic chip with a novel isa and on-chip learning. *National Science Review*, page nwae102, 2024.
- [Merolla *et al.*, 2014] Paul A Merolla, John V Arthur, Rodrigo Alvarez-Icaza, Andrew S Cassidy, Jun Sawada, Filipp Akopyan, et al. A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science*, 345(6197):668–673, 2014.
- [Na *et al.*, 2022] Byunggook Na, Jisoo Mok, Seongsik Park, Dongjin Lee, Hyeokjun Choe, and Sungroh Yoon. Autosnn: Towards energy-efficient spiking neural networks. In *International Conference on Machine Learning*, pages 16253–16269. PMLR, 2022.
- [Orchard *et al.*, 2021] Garrick Orchard, E Paxon Frady, Daniel Ben Dayan Rubin, Sophia Sanborn, Sumit Bam Shrestha, Friedrich T Sommer, and Mike Davies. Efficient neuromorphic signal processing with loihi 2. In *2021 IEEE Workshop on Signal Processing Systems*, pages 254–259. IEEE, 2021.
- [Rathi and Roy, 2021] Nitin Rathi and Kaushik Roy. 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 *et al.*, 2018] Nitin Rathi, Priyadarshini Panda, and Kaushik Roy. Stdp-based pruning of connections and weight quantization in spiking neural networks for energy-efficient recognition. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 38(4):668–677, 2018.
- [Rueckauer *et al.*, 2017] Bodo Rueckauer, Iulia-Alexandra Lungu, Yuhuang Hu, Michael Pfeiffer, and Shih-Chii Liu. Conversion of continuous-valued deep networks to efficient event-driven networks for image classification. *Frontiers in Neuroscience*, 11:682, 2017.
- [Sehwag *et al.*, 2020] Vikash Sehwag, Shiqi Wang, Prateek Mittal, and Suman Jana. Hydra: Pruning adversarially robust neural networks. *Advances in Neural Information Processing Systems*, 33:19655–19666, 2020.
- [Shen *et al.*, 2023] Jiangrong Shen, Qi Xu, Jian K Liu, Yueming Wang, Gang Pan, and Huajin Tang. Esl-snns: An evolutionary structure learning strategy for spiking neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 86–93, 2023.
- [Srinivasan and Roy, 2019] Gopalakrishnan Srinivasan and Kaushik Roy. Restocnet: Residual stochastic binary convolutional spiking neural network for memory-efficient neuromorphic computing. *Frontiers in Neuroscience*, 13:189, 2019.
- [Wang *et al.*, 2020] Tianzhe Wang, Kuan Wang, Han Cai, Ji Lin, Zhijian Liu, Hanrui Wang, Yujun Lin, and Song Han. Apq: Joint search for network architecture, pruning and quantization policy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2078–2087, 2020.
- [Wang *et al.*, 2023] Yang Wang, Bo Dong, Yuji Zhang, Yunduo Zhou, Haiyang Mei, Ziqi Wei, and Xin Yang. Event-enhanced multi-modal spiking neural network for dynamic obstacle avoidance. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 3138–3148, 2023.
- [Warden, 2018] Pete Warden. Speech commands: A dataset for limited-vocabulary speech recognition. *arXiv preprint arXiv:1804.03209*, 2018.
- [Wei *et al.*, 2024] Wenjie Wei, Malu Zhang, Jilin Zhang, Ammar Belatreche, Jibin Wu, Zijing Xu, Xuerui Qiu, Hong Chen, Yang Yang, and Haizhou Li. Event-driven learning for spiking neural networks. *arXiv preprint arXiv:2403.00270*, 2024.
- [Wu *et al.*, 2021] Jibin Wu, Yansong Chua, Malu Zhang, Guoqi Li, Haizhou Li, and Kay Chen Tan. A tandem learning rule for effective training and rapid inference of deep spiking neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 2021.
- [Xu *et al.*, 2023] Qi Xu, Yaxin Li, Jiangrong Shen, Jian K Liu, Huajin Tang, and Gang Pan. Constructing deep spiking neural networks from artificial neural networks with knowledge distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7886–7895, 2023.
- [Yan *et al.*, 2024] Jiaqi Yan, Qianhui Liu, Malu Zhang, Lang Feng, De Ma, Haizhou Li, and Gang Pan. Efficient spiking neural network design via neural architecture search. *Neural Networks*, page 106172, 2024.
- [Yang *et al.*, 2022] Qu Yang, Qi Liu, and Haizhou Li. Deep residual spiking neural network for keyword spotting in low-resource settings. In *INTERSPEECH*, pages 3023–3027, 2022.
- [Yılmaz *et al.*, 2020] Emre Yılmaz, Özgür Bora Gevrek, Jibin Wu, Yuxiang Chen, Xuanbo Meng, and Haizhou Li. Deep convolutional spiking neural networks for keyword spotting. In *INTERSPEECH*, pages 2557–2561, 2020.
- [Zhang *et al.*, 2016] Shijin Zhang, Zidong Du, Lei Zhang, Huiying Lan, Shaoli Liu, Ling Li, Qi Guo, Tianshi Chen, and Yunji Chen. Cambricon-x: An accelerator for sparse neural networks. In *2016 49th Annual IEEE/ACM International Symposium on Microarchitecture*, pages 1–12. IEEE, 2016.
- [Zhang *et al.*, 2021] Malu Zhang, Jiadong Wang, Jibin Wu, Ammar Belatreche, Burin Amornpaisannon, Zhixuan Zhang, et al. Rectified linear postsynaptic potential function for backpropagation in deep spiking neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 33(5):1947–1958, 2021.
- [Zheng *et al.*, 2021] Hanle Zheng, Yujie Wu, Lei Deng, Yifan Hu, and Guoqi Li. Going deeper with directly-trained larger spiking neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 11062–11070, 2021.