This CVPR paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Spiking Transformer with Spatial-Temporal Attention

Donghyun Lee, Yuhang Li, Youngeun Kim, Shiting Xiao, Priyadarshini Panda Department of Electrical Engineering, Yale University

{donghyun.lee, yuhang.li, youngeun.kim, ginny.xiao, priya.panda}@yale.edu

Abstract

Spike-based Transformer presents a compelling and energyefficient alternative to traditional Artificial Neural Network (ANN)-based Transformers, achieving impressive results through sparse binary computations. However, existing spike-based transformers predominantly focus on spatial attention while neglecting crucial temporal dependencies inherent in spike-based processing, leading to suboptimal feature representation and limited performance. To address this limitation, we propose Spiking Transformer with Spatial-Temporal Attention (STAtten), a simple and straightforward architecture that efficiently integrates both spatial and temporal information in the selfattention mechanism. STAtten introduces a block-wise computation strategy that processes information in spatialtemporal chunks, enabling comprehensive feature capture while maintaining the same computational complexity as previous spatial-only approaches. Our method can be seamlessly integrated into existing spike-based transformers without architectural overhaul. Extensive experiments demonstrate that STAtten significantly improves the performance of existing spike-based transformers across both static and neuromorphic datasets, including CIFAR10/100, ImageNet, CIFAR10-DVS, and N-Caltech101. The code is available at https://github.com/Intelligent-Computing-Lab-Yale/STAtten.

1. Introduction

Spiking Neural Networks (SNNs) are emerging as a promising alternative to conventional Artificial Neural Networks (ANNs) [35, 40] due to their bio-inspired, energy-efficient computing paradigm. By leveraging sparse binary spikebased computations, SNNs can achieve significant energy savings while enabling deployment across various neuromorphic computing platforms such as TrueNorth [1], Loihi [6], and Tianjic [38]. However, the challenge lies in effectively processing and learning from these discrete, timedependent spike patterns while maintaining computational efficiency and accuracy. Traditional SNNs have primarily



Figure 1. Heatmaps of spatial-only attention versus our STAtten on sequential CIFAR100 dataset. Input images are divided column-wise, where each column corresponds to one timestep.

relied on convolution-based architectures adapted from successful ANN models like VGGNet [41, 49] and ResNet [20, 60]. While these architectures benefit from computational efficiency, they often struggle with binary spike activations, leading to information loss and degraded performance. This limitation has driven researchers to explore alternative architectures that can better handle sparse, temporal data while maintaining the energy benefits of spike-based computation. Transformer architectures, with their remarkable success across various domains [2, 4, 11, 47], have emerged as a promising direction for enhancing SNN capabilities. Recent works have introduced spike-formed transformers that adapt the transformer's self-attention mechanism to the spike-based domain for diverse tasks, such as, object tracking [58], computer vision [42, 50, 52, 53, 62, 64], depth estimation [59], and speech recognition [45, 48]. These architectures, such as Spikformer [64] and Spike-driven Transformer [53], implement innovative binary representations for query (O), key (K), and value (V) computations while eliminating the computationally expensive Softmax function. This binary attention design preserves the essential characteristics of spike-based computation while enabling more efficient information processing compared to conventional transformer architectures. However, current spike-formed transformers predominantly rely on spatialonly attention, overlooking the inherently dynamic and temporal nature of spike events. To understand these different attention patterns, we compare spatial-only and spatialtemporal attention on sequential CIFAR100, as shown in Fig. 1. The visualization reveals that spatial-only attention focuses solely on vertical (spatial) relationships within each timestep, missing crucial object features that evolve over time. In contrast, our spatial-temporal attention captures both vertical (spatial) and horizontal (temporal) relationships, enabling comprehensive feature representation of sequential inputs.

In this work, we introduce Spiking Transformer with Spatial-Temporal Attention (STAtten), which integrates both spatial and temporal information within the selfattention mechanism. Through analysis of memory requirements and neuronal activity patterns, we design a blockwise computation strategy with local temporal correlations. By eliminating Softmax operation, we maintain the same computational complexity ($O(TND^2)$) as existing spikebased transformers [42, 52, 62, 64]. Furthermore, our approach can be flexibly integrated into existing architectures without structural modifications, enhancing their representational capacity. Experiments on both static and dynamic datasets demonstrate consistent performance improvements over spatial-only attention architectures. The main contributions of our work are as follows:

- We identify limitations in current spiking transformers that rely on spatial-only attention through empirical analysis, demonstrating the importance of capturing temporal dependencies inherent in spike-based processing.
- We propose STAtten, a block-wise spatial-temporal attention mechanism for spike-based transformers that maintains the original computational complexity $(O(TND^2))$.
- We introduce a flexible plug-and-play design that enables STAtten to be integrated into existing spike-based transformers without compromising their core architecture.
- Through extensive experiments across both static (CI-FAR10/100, ImageNet) and neuromorphic (CIFAR10-DVS, N-Caltech101) datasets, we demonstrate that STAtten consistently improves performance across different backbone architectures.

2. Related Works

Training Strategies for SNNs. Significant progress has been made in developing effective training strategies for SNNs, primarily following two approaches: ANN-to-SNN conversion and direct training. The conversion approach transforms pre-trained ANNs into SNNs, leveraging established ANN architectures to enhance SNN performance. Several studies [8, 10, 17, 30] have formalized this conversion process and demonstrated its effectiveness in transferring ANN knowledge to the spike domain. Direct training through backpropagation, while fundamental to ANNs, presents unique challenges in SNNs due to the

non-differentiable nature of spike events. The introduction of surrogate gradients has been crucial in enabling effective backpropagation-based training in SNNs, as demonstrated in pioneering works [5, 43, 49]. Notably, direct training with surrogate gradients has shown superior performance compared to conversion methods across various tasks, including image classification [14, 20, 31, 55, 57], semantic segmentation [23], and object detection [44], while maintaining the energy efficiency inherent to spike-based computation.

Spiking Transformer. Despite the energy efficiency of convolution-based SNNs, their performance often lags behind ANNs due to information loss from sparse binary activations. To address this limitation, researchers have adapted transformer architectures to the SNN domain, leveraging their powerful self-attention mechanisms and ability to capture long-range dependencies. This adaptation has led to several innovative spike-formed transformer architectures. Spikformer [64] introduced the first SNN-based transformer architecture, establishing foundational principles for spikeformed self-attention mechanisms. Their key innovation lies in eliminating the Softmax function from self-attention computations, justified by two key observations: (1) the binary nature of spike-based Q, K, and V values, and (2) the redundancy of Softmax scaling for binarized operations. Building upon this foundation, Spike-driven Transformer [53] enhanced the architecture with a novel selfattention block incorporating masking and sparse addition to optimize power efficiency. Additionally, inspired by MS-ResNet [20], Spike-driven Transformer reimagined residual connections to propagate membrane potential rather than spikes. Recent works have further extended these concepts, leading to architectures like SpikingReformer [42], QK-Former [62], and Spike-driven Transformer-V2 [52], each introducing unique optimizations for spike-based processing. However, existing spiking transformer architectures primarily focus on spatial attention, neglecting the temporal dynamics inherent in spike-based computation.

3. Preliminary

Leaky Integrate-and-Fire Neuron. As a fundamental building block of SNNs, the Leaky Integrate-and-Fire (LIF) neuron [3] has emerged as an important component for energy-efficient computation. The LIF neuron is a nonlinear activation function that determines whether neurons fire spikes as follows:

$$\mathbf{u}[t+1]^{l} = \tau \mathbf{u}[t]^{l} + \mathbf{W}^{l} f(\mathbf{u}[t]^{l-1})$$
(1)

$$f(\mathbf{u}[t]^{l}) = \begin{cases} 1 & \text{if } \mathbf{u}[t]^{l} > V_{th}, \\ 0 & \text{otherwise} \end{cases},$$
(2)

where, $\mathbf{u}[t]^l$ is the membrane potential in *l*-th layer at timestep $t, \tau \in (0, 1]$ is the leaky factor for membrane po-



Figure 2. Comparison between different self-attentions with the CIFAR100 dataset. Analysis of entropy and accuracy for Temporal-only (T), Spatial-only (S), and Spatial-temporal (ST).

tential leakage, \mathbf{W}^{l} is the weight of *l*-th layer, and $f(\cdot)$ is the LIF function with firing threshold V_{th} . Therefore, when the membrane $\mathbf{u}[t]^{l}$ is higher than V_{th} , the LIF function fires a spike and the membrane potential is reset to 0.

Vanilla Self-attention. Transformers [11, 47] have surpassed convolution-based architectures due to their unique self-attention mechanism, which can capture global features through long-range dependency. For a floating-point input tensor $\mathbf{X}_f \in \mathbb{R}^{N \times D}$ with N tokens, and D features, we formulate the self-attention as follows:

$$\mathbf{Q}_f = \mathbf{W}_{\mathbf{Q}} \mathbf{X}_f, \quad \mathbf{K}_f = \mathbf{W}_f \mathbf{X}_f, \quad \mathbf{V}_f = \mathbf{W}_{\mathbf{V}} \mathbf{X}_f \quad (3)$$

Attn = Softmax
$$(\frac{\mathbf{Q}\mathbf{K}^{\top}}{\sqrt{D}})\mathbf{V}, \quad \mathbf{Q}_f, \mathbf{K}_f, \mathbf{V}_f \in \mathbb{R}^{N \times D},$$
(4)

where $\mathbf{W}_{\mathbf{Q}}$, $\mathbf{W}_{\mathbf{K}}$, $\mathbf{W}_{\mathbf{V}} \in \mathbb{R}^{D \times D}$ are linear projections for floating-point \mathbf{Q}_f , \mathbf{K}_f , \mathbf{V}_f , respectively. Note that the computational complexity of attention is $\mathcal{O}(N^2D)$ due to its matrix multiplication operation.

Spike-based Self-attention. The first spike-based transformer work, Spikformer [64], proposes Spiking Self Attention (SSA) with binary \mathbf{Q} , \mathbf{K} , \mathbf{V} . For binary input tensor $\mathbf{X} \in \mathbb{R}^{T \times N \times D}$ with T timestep, we formulate the SSA as follows:

$$\mathbf{Q}, \mathbf{K}, \mathbf{V} = \text{LIF}(\mathbf{W}_{\mathbf{Q}}\mathbf{X}), \text{LIF}(\mathbf{W}_{\mathbf{K}}\mathbf{X}), \text{LIF}(\mathbf{W}_{\mathbf{K}}\mathbf{X}),$$
$$\text{SSA}(\mathbf{Q}\mathbf{K}\mathbf{V}) = \text{LIF}(\mathbf{Q}\mathbf{K}^{\top}\mathbf{V} \cdot \alpha)$$
(5)

where $\mathbf{W}_{\mathbf{Q}}, \mathbf{W}_{\mathbf{K}}, \mathbf{W}_{\mathbf{V}} \in \mathbb{R}^{D \times D}$ are linear projections for binary $\mathbf{Q}, \mathbf{K}, \mathbf{V}$, respectively, and α denotes a scaling factor.

4. Methodology

In this section, we present our analysis of spike-based transformers and provide the details of our proposed methodology. We first establish the necessity of spatial-temporal attention in spike-based transformers through an informationtheoretic analysis based on entropy measurements. Building upon these observations, we introduce our STAtten mechanism.



Figure 3. (a) Maximum batch size of block-wise STAtten and full spatial-temporal attention (without block partitioning) for running on A5000 GPU with 24GB VRAM memory. (b) Average number of active neurons after QKV computation at different timestep combinations, where [t] indicates timestep index.

4.1. Motivation of Spatial-Temporal Attention

The feature maps of spike-based transformers at each timestep carry distinct information, as their activation distributions are dynamically influenced by both the leaky factor and reset mechanisms. This temporal aspect of information flow raises a fundamental question: how do different attention mechanisms effectively capture and utilize this spatiotemporal information?

First, we define the three different types of attention: Spatial, Temporal, and Spatial-temporal attention with timestep t, token position n, and feature dimension d as follows:

$$\begin{split} \mathbf{S}_{-} \mathrm{Attn}[t,:,:] &= \mathrm{LIF}(\mathbf{Q}[t,:,:]\mathbf{K}^{\top}[t,:,:]\mathbf{V}[t,:,:] \cdot \alpha), \\ \mathbf{T}_{-} \mathrm{Attn}[:,n,:] &= \mathrm{LIF}(\mathbf{Q}[:,n,:]\mathbf{K}^{\top}[:,n,:]\mathbf{V}[:,n,:] \cdot \alpha), \\ \mathrm{ST}_{-} \mathrm{Attn}[:,:,d] &= \mathrm{LIF}(\mathbf{Q}[:,:,d]\mathbf{K}^{\top}[:,:,d]\mathbf{V}[:,:,d] \cdot \alpha), \end{split}$$

where S_Attn, T_Attn, and ST_Attn are spatial, temporal, and spatial-temporal attention, and the **Q**, **K**, **V** dimensions of S_Attn, T_Attn, and ST_Attn are $\mathbb{R}^{T \times N \times D}$. The key distinction lies in their attention computation scope: spatial attention operates on token relationships within individual timesteps, temporal attention focuses on cross-timestep feature dependencies, and spatial-temporal attention integrates both spatial and temporal feature correlations. To quantify the information distribution in each attention mechanism, we measure Shannon entropy [16] as follows:

$$H(\text{Attn}) = -\sum_{t=1}^{T} \sum_{n=1}^{N} \sum_{d=1}^{D} \hat{\mathbf{p}}[t, n, d] \log \hat{\mathbf{p}}[t, n, d], \quad (7)$$

where $\hat{\mathbf{p}}[t, n, d] = \text{Softmax}(\text{Attn}[t, n, d])$. Lower entropy values indicate more focused spike patterns, while higher entropy values suggest dispersed patterns [25]. We analyze this relationship using a pretrained Spike-driven Transformer [53] on the CIFAR100 dataset, as shown in Fig. 2. The results reveal an inverse correlation between entropy and accuracy: temporal-only attention shows high entropy (5.81) but lowest accuracy (77.7%), while



Figure 4. Overview of STAtten architecture. (a) Block-wise temporal attention mechanism. Binary Q, K, V tensors are partitioned into temporal blocks, where black lines indicate paired timestep processing. (b) Computation flow with tensor dimensions, where T is the number of timesteps, N is the number of tokens, B is block size, and D is the feature dimension.

spatial-temporal attention achieves lower entropy (4.85) with higher accuracy (79.9%). These findings verify that combining spatial and temporal processing leads to more structured feature representations, providing strong motivation for spatial-temporal attention.

However, we argue that fully correlating temporal information across all timesteps like ST_Attn in Eq. 6 introduces two challenges: (1) substantial memory requirements for storing temporal information, and (2) increased dead neurons when correlating distant timesteps due to weak spike similarity. The first challenge arises from matrix multiplication with large temporal dimensions (TN, D). To address this, we partition the temporal dimension into blocks, reducing the memory footprint for matrix operations. As shown in Fig. 3(a), this block-wise computation strategy leads to around $1.6 \times$ higher batch size compared to full temporal attention. The second challenge stems from binary matrix multiplication in spike-based processing. Our analysis in Fig. 3(b) reveals that QKV computations show higher active neurons when processing same-timestep, while activity decreases with temporal distance.

4.2. STAtten Mechanism

Based on these observations, we propose Spatial-Temporal Attention (STAtten). The block-wise STAtten partitions the temporal sequence into blocks of size B, as shown in Fig. 4(a). We define STAtten as follows:

STAtten($\mathbf{X}[b]$) = LIF($\mathbf{Q}[b]\mathbf{K}^{\top}[b]\mathbf{V}[b] \cdot \alpha$), where $[b] = [iB : (i+1)B, :, d], \quad i \in \{0, 1, ..., T/B - 1\}$ (8)

where [iB : (i + 1)B, :, d] denotes the aggregated features from timesteps iB to (i + 1)B. For instance, with T = 4and B = 2, features are grouped in pairs of timesteps [0,1], [2,3]. By processing blocks of timesteps together, we capture local temporal relationships while maintaining manageable computation. This approach requires storing only a subset of temporal information at any given time, reduc-

Table 1. Computational complexity and energy consumption of different attention mechanisms. E_{MAC} and E_{AC} are energy costs for MAC and AC operations, N, D, T denote patches, dimension, timesteps, and S_Q , S_K , S_V are firing rates of Q, K, V. SDT represents Spike-driven Transformer.

Method	ST	Complexity	Energy
ViT [11] ViViT [2]	× ✓	$ \begin{array}{c} \mathcal{O}(N^2 D) \\ \mathcal{O}(T^2 N^2 D) \end{array} $	$\frac{E_{MAC} \cdot N^2 D}{E_{MAC} \cdot T^2 N^2 D}$
Spikformer [64] SDT [53] SDT-V2 [52]	× × ×	$\mathcal{O}(TND^2)$ $\mathcal{O}(TND)$ $\mathcal{O}(TND^2)$ $\mathcal{O}(TND^2)$	$E_{AC} \cdot TND^2 \cdot (S_Q + S_K + S_V)$ $E_{AC} \cdot TND \cdot (S_Q + S_K)$ $E_{AC} \cdot TND^2 \cdot (S_Q + S_K + S_V)$ $E_{AC} \cdot TND^2 \cdot (S_Q + S_K + S_V)$
STAtten	~~	$\mathcal{O}(TND^2)$ $\mathcal{O}(TND^2)$	$\frac{E_{AC} \cdot TND^2 \cdot (S_Q + S_K + S_V)}{E_{AC} \cdot TND^2 \cdot (S_Q + S_K + S_V)}$

ing memory overhead. We set block size B = 2 with 4 timesteps for static datasets and B = 4 with 16 timesteps for neuromorphic datasets.

Unlike traditional transformers that require Softmax operations, the binary \mathbf{Q} , \mathbf{K} , and \mathbf{V} allow flexible reordering of computations. Generalizing to arbitrary number of timesteps T, the computation follows:

$$\mathbf{Q}, \mathbf{K}, \mathbf{V} \in \mathbb{R}^{\frac{T}{B} \times N \times D} \xrightarrow{\text{reshape}} \mathbf{Q}, \mathbf{K}, \mathbf{V} \in \mathbb{R}^{(\frac{T}{B}N) \times D}$$
$$\mathbf{K}^{\top} \mathbf{V} \in \mathbb{R}^{D \times D} \xrightarrow{\text{attention}} \mathbf{Q}(\mathbf{K}^{\top} \mathbf{V}) \in \mathbb{R}^{\frac{T}{B} \times N \times D}.$$
(9)

The computation sequence reduces intermediate memory requirements through early $\mathbf{K}^{\top}\mathbf{V}$ calculation while preserving cross-timestep information flow within blocks. This implementation maintains computational complexity of $\mathcal{O}(TND^2)$, matching spatial-only attention [42, 52, 53, 62, 64] while incorporating temporal dependencies. By avoiding the quadratic complexity in timesteps and the number of tokens (T^2N^2) that would result from full temporal correlation, our method scales efficiently to longer sequences. Fig. 4(b) illustrates the complete computation process with corresponding tensor dimensions. Table 2. Performance comparison on sequential CIFAR10/100 datasets, demonstrating the effectiveness of spatial-temporal attention in capturing long-term dependencies.

Method	s-CIFAR10 (%)	s-CIFAR100 (%)
PLIF [14]	81.47	53.38
Spikformer [64]	79.29	57.17
STAtten + [64]	82.45	58.75
Spike-driven Transformer [53]	80.32	61.08
STAtten + [53]	83.41	64.30
SpikingReformer [42]	79.25	57.48
STAtten + [42]	81.84	58.30

4.3. STAtten with Existing Spiking Transformers

STAtten is designed to be a plug-and-play module that can enhance various spike-based transformer architectures, including Spikformer [64], Spike-driven Transformer [53], Spike-driven Transformer-V2 [52], SpikingReformer [42], and QKformer [62] while preserving their unique characteristics.

The key to successful integration lies in preserving each architecture's distinct residual connection strategies while replacing their spatial-only attentions with STAtten. Two primary residual connection mechanisms are prevalent in current architectures: Spike Element-Wise (SEW) shortcut [13] and Membrane-Shortcut (MS) [20]. In architectures employing SEW shortcuts, such as Spikformer and QK-Former, spike activations are directly added to the attention output. In contrast, Spike-driven Transformer, Spike-driven Transformer-V2, and SpikingReformer utilize MS-shortcuts that propagate membrane potentials rather than spikes Our STAtten module maintains compatibility with both connection types. We detail the existing spike-based attention mechanisms of each architecture and our integration strategy.

Spikformer [64]: Spikformer proposes Spiking Self Attention (SSA) module which is inspired by the vanilla self-attention of ViT [11] as follows:

$$SSA(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{LIF}(\mathbf{Q} \odot \mathbf{K}^{\top} \odot \mathbf{V} \cdot \alpha), \quad (10)$$

where \odot is matrix multiplication. This work establishes two key principles for spike-based attention: non-softmax operation and flexible **Q**, **K**, **V** operation ordering. While these principles form the foundation for subsequent spike-based architectures, the SSA module processes features independently at each timestep. We address this limitation by replacing the attention mechanism with STAtten.

Spike-driven Transformer [53]: Spike-driven Transformer introduces Spike-Driven Self-Attention (SDSA), which reformulates attention computation as:

$$SDSA(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \mathbf{Q} \otimes LIF(\sum_{c} (\mathbf{K} \otimes \mathbf{V})),$$
 (11)

where \otimes denotes element-wise multiplication and \sum_c represents column-wise summation. While SDSA achieves computational efficiency through element-wise operations, this design limits feature correlation even within the spatial domain due to the element-wise multiplication. By integrating STAtten into this architecture, we enable comprehensive spatial-temporal feature correlation.

SpikingReformer [61]: SpikingReformer identifies two key challenges in adapting conventional **Q**, **K**, **V** attention to spike-based transformers: (1) floating-point matrix multiplication between **Q**, **K**, **V** and (2) Softmax operations involving exponentiation and division. To address these challenges, the authors propose Dual Spike Self-Attention (DSSA):

$$DSSA(\mathbf{X}) = \text{LIF}(\mathbf{X} \odot \mathbf{W}^{\top} \odot \mathbf{X}^{\top} \cdot \alpha), \quad (12)$$

where X and W denote binary input and linear projection, respectively. We try to verify that these concerns can be addressed by replacing DSSA with STAtten, showing that Q, K, V matrix multiplication is still compatible with spikebased processing when properly designed.

QKFormer [62]: QKFormer implements dual attention mechanisms: Q-K Channel Attention (QKTA) and SSA. While SSA follows the Spikformer design, QKTA introduces a token-wise attention:

$$QKTA = \text{LIF}(\sum_{N} (\mathbf{Q})) \otimes \mathbf{K}, \tag{13}$$

where \sum_{N} represents token-wise summation. QKFormer redesigns the architecture by incorporating embedding layers in each encoder. In this architecture, we only replace SSA with STAtten as SSA is the primary mechanism responsible for capturing feature relationships.

As we only modify the self-attention module, the original information flow through residual connections remains unchanged, maintaining the architecture's core characteristics. After computing STAtten across temporal blocks, the features are concatenated and processed through subsequent layers: $\mathbf{F}_{out} = \text{Concat}(\text{STAtten}(\mathbf{X}[iB : (i+1)B])_{i=0}^{T/B-1})$. This concatenated feature then flows through convolutional and batch normalization layers: $\mathbf{Z} = \text{BN}(\text{Conv}(\mathbf{F}_{out}))$.

4.4. Complexity and Energy of Self-attention

We analyze both computational complexity and theoretical energy consumption of our proposed architecture, focusing specifically on the self-attention mechanism. A comprehensive evaluation of the overall architecture's energy profile is provided in Supplementary Material.

As shown in Table 1, despite incorporating spatialtemporal attention capabilities, our STAtten maintains the same computational complexity $\mathcal{O}(TND^2)$ as previous spatial-only spiking transformers. This efficiency is

Table 3. Performance comparison between our methods and previous works on CIFAR10 and CIFAR100 datasets. In the architecture column, Architecture-L-D represents L number of encoder blocks and D hidden dimensions. All baseline architectures of spike-based transformer and corresponding STAtten implementations are trained from scratch for fair comparison.

Method	Architecture	Timestep	CIFAR10 (%)	CIFAR100 (%)
tdBN [60]	ResNet19	4	92.92	-
PLIF [14]	ConvNet	8	93.50	-
Dspike [31]	ResNet18	6	94.25	74.24
DSR [36]	ResNet18	20	95.40	78.50
DIET-SNN [39]	VGG16	5	92.70	69.67
SNASNet [24]	ConvNet	5	93.64	73.04
Spikformer [64]	Spikformer-4-384	4	93.99	75.06
STAtten + [64]	Spikformer-4-384	4	94.36	75.85
Spike-driven Transformer [53]	Spike-driven Transformer-2-512	4	95.60	78.40
STAtten + [53]	Spike-driven Transformer-2-512	4	96.03	79.85
SpikingReformer [42]	SpikingReformer-6-384	4	95.03	77.16
STAtten + [42]	SpikingReformer-6-384	4	95.26	77.90
QKFormer [62]	QKFormer-4-384	4	95.12	79.79
STAtten + [62]	QKFormer-4-384	4	95.35	80.20

achieved through our non-Softmax design, which enables flexible reordering of Q, K, V operations. In contrast, conventional ANN-based spatial-temporal attention methods like ViViT [2] require substantially higher complexity ($\mathcal{O}(T^2N^2D)$) due to their softmax-constrained computation order. For energy consumption estimation, we follow the floating-point operations (FLOPs) implementation in 45nm technology [18]. The energy costs are categorized into E_{MAC} for multiply-accumulate operations corresponding to 32-/8-bit ANN computations, and E_{AC} for accumulate operations used in binary SNN calculations. The final energy consumption is modulated by the firing rates (S_Q , S_K , and S_V) of the respective Q, K, and V. A detailed analysis of these energy measurements and their implications is presented in Section 5.3.

5. Experiments

We verify our proposed methods on both static and dynamic datasets, *e.g.*, CIFAR10/100 [26], ImageNet [7], CIFAR10-DVS [29], and N-Caltech101 [37] datasets. We use direct coding to convert float pixel values into binary spikes [49]. In all our experiments, the spiking transformer STAtten models are trained from scratch. The detailed training strategy for each dataset is explained in the Supplementary Material.

Notation: In the experiments, we use the notation STAtten $+ [\cdot]$ to denote STAtten implementation based on the backbone architecture used in $[\cdot]$. E.g., STAtten+[64] refers to training a STAtten model with Spikformer backbone from scratch.

5.1. Sequential CIFAR10/100 Classification

We first evaluate our spatial-temporal attention mechanism's capability to capture long-term temporal dependencies using sequential CIFAR10 and CIFAR100 datasets. In this task, we transform the standard image classification task into a temporal sequence processing task by dividing each input image column-wise, shown in Fig. 1, where each column serves as input for one timestep. This results in a sequence length of 32 timesteps, corresponding to the image width. To accommodate this 1-D sequential input, we modify the convolutional, batch normalization, and pooling layers from 2-D to 1-D operations. As shown in Table 2, STAtten outperforms existing spatial-only attention methods on both datasets. Specifically, in Spike-driven Transformer [53], our approach achieves improvements of 3.09 % and 3.22 % on s-CIFAR10 and s-CIFAR100, respectively, compared to the baseline models. This performance gain demonstrates STAtten's capability in capturing long-range temporal dependencies.

5.2. Performance Analysis

CIFAR10/100. The CIFAR10/100 datasets comprise static images of 32×32 pixels, with 50,000 training and 10,000 test images. As shown in Table 3, STAtten improves the performance of various spike-based transformers, achieving 95.35% accuracy on CIFAR10 and 80.20% on CIFAR100 when integrated with QKFormer. These improvements over spatial-only attention methods are consistent across different architectures, demonstrating the benefits of spatial-temporal processing.

ImageNet. ImageNet is a large-scale image dataset with images sized 224×224 pixels, comprising approximately 1.2 million training images and 50,000 validation images across 1,000 classes. As shown in Table 4, STAtten enhances the performance of spike-based transformers. With Spike-driven Transformer [53], it achieves 78.11% accuracy at 224 × 224 resolution and 78.39% at 288 × 288 resolution. Integration with Spike-driven Transformer-V2

Table 4. Performance comparison between our methods and previous works on ImageNet datasets. Note that * represents the inference accuracy with 288×288 resolution. In the architecture column, *L*-*D* represents *L* number of encoder blocks and *D* hidden dimensions. \dagger represents our implementation.

Method	Туре	Architecture	Param(M)	Timestep	Accuracy(%)
Vision Transformer [11]	ANN	ViT-B/16	86.59	1	77.90
DeiT [46]	ANN	DeiT-B	86.59	1	81.80
Swin Transformer [33]	ANN	Swin Transformer-B	87.77	1	83.50
TET [9]	SNN	SEW-ResNet34	21.79	4	68.00
Spiking ResNet [19]	SNN	ResNet50	25.56	350	72.75
tdBN [60]	SNN	ResNet34	21.79	6	63.72
SEW-ResNet [13]	SNN	SEW-ResNet152	60.19	4	69.26
MS-ResNet [20]	SNN	ResNet104	78.37	5	74.21 / 76.02*
Att-MS-ResNet [51]	SNN	ResNet104	78.37	4	77.08*
Spikformer [64]	SNN	Spikformer-8-768	66.34	4	74.81
SpikingReformer [42]	SNN	SpikingReformer-6-1024	66.38	4	78.77 / 79.40*
QKFormer [62]	SNN	QKFormer-6-1024	64.96	4	84.22
	SNN	Spike-driven Transformer-8-512	29.68	4	74.57
Spike-driven Transformer [53]	SNN	Spike-driven Transformer-10-512	36.01	4	74.66
	SNN	Spike-driven Transformer-8-768	66.34	4	76.32 / 77.07*
STAtton + [52]	SNN	Spike-driven Transformer-8-512	29.68	4	76.18 / 76.56*
STAtten + [55]	SNN	Spike-driven Transformer-8-768	66.34	4	78.11 / 78.39*
Spike-driven Transformer-V2 [52]	SNN	Spike-driven Transformer-V2-8-512	55.4	4	79.49† / 79.98*†
STAtten + [52]	SNN	Spike-driven Transformer-V2-8-512	55.4	4	79.85 / 80.67*

Table 5. Performance comparison between our methods and previous works on CIFAR10-DVS and N-Caltech101 datasets. In the architecture column, L-D represents L number of encoder blocks and D hidden dimensions. All baseline architectures of spike-based transformer and corresponding STAtten implementations are trained from scratch for fair comparison.

Method	Architecture	Timestep	CIFAR10-DVS (%)	N-Caltech101 (%)
tdBN [60]	ResNet19	10	67.80	-
PLIF [14]	ConvNet	20	74.80	-
Dspike [31]	ResNet18	10	75.40	-
DSR [36]	ResNet18	20	77.27	-
SEW-ResNet [13]	ConvNet	20	74.80	-
TT-SNN [27]	ResNet34	6	-	77.80
NDA [32]	VGG11	10	79.6	78.2
Spikformer [64]	Spikformer-2-256	16	80.9	-
Spike-driven Transformer [53]	Spike-driven Transformer-2-256	16	80.0	81.80
STAtten + [53]	Spike-driven Transformer-2-256	16	81.1	83.15
SpikingReformer [42]	SpikingReformer-4-384	16	78.80	81.29
STAtten + [42]	SpikingReformer-4-384	16	80.60	81.95
QKFormer [62]	_QKFormer-4-384	16	82.90	83.58
STAtten + [62]	SpikingReformer-4-384	16	83.90	84.25

[52] further improves performance to 79.85% and 80.67% at respective resolutions. It is noteworthy that Spike-driven Transformer [53] achieves comparable or better performance than recent Spikformer models while using only 29.68M parameters. Furthermore, STAtten with Spike-driven Transformer-V2 achieves even higher accuracy than SpikingReformer and Spikformer while using around 20% fewer parameters.

Neuromorphic Datasets. We evaluate our approach on two neuromorphic vision datasets: CIFAR10-DVS [29] and N-Caltech101 [37]. CIFAR10-DVS contains 10,000 sam-

ples (9,000 training, 1,000 testing) captured using a Dynamic Vision Sensor (DVS) from static CIFAR10 images. N-Caltech101, derived from the Caltech101 dataset [28], comprises 8,831 DVS image samples. Following standard protocols, we resize CIFAR10-DVS samples to 64×64 and apply the augmentation strategy from [53], while N-Caltech101 samples are resized to 64×64 with NDA [32] augmentation. As shown in Table 5, integrating STAtten consistently improves performance across different architectures. With Spike-driven Transformer as the backbone, our approach achieves 81.1% accuracy on CIFAR10-DVS

Table 6. Comparison of memory, energy, and accuracy on ImageNet across different transformer architectures. The precision format weight/activation bit-width and the memory include the size of model parameters and activations. We show the memory consumption for both 32-bit and 8-bit weight spiking models. SDT represents Spike-driven Transformer.

Precision	Timestep	Memory (MB)	Energy (mJ)
32/32	1	351.8	254.84
8/8	1	87.96	110.8
32(8)/1	4	285.99 (86.48)	21.48
32(8)/1	4	273.57 (85.4)	8.76
32(8)/1	4	280.28 (84.99)	38.91
32(8)/1	4	283.97 (87.42)	12.42
32(8)/1	4	283.97 (87.42)	12.36
32(8)/1	4	250.22 (84.02)	52.40
32(8)/1	4	250.22 (84.02)	52.38
	Precision - 32/32 - 8/8 32(8)/1 32(8)/1 - 32(8)/1 - 32(8)/1 - 32(8)/1 - 32(8)/1 32(8)/1 32(8)/1 32(8)/1	$\begin{array}{c c c c c c c c c c c c c c c c c c c $	Precision Timestep Memory (MB) 32/32 1 351.8 8/8 1 87.96 32(8)/1 4 285.99 (86.48) 32(8)/1 4 273.57 (85.4) 32(8)/1 4 280.28 (84.99) 32(8)/1 4 283.97 (87.42) 32(8)/1 4 283.97 (87.42) 32(8)/1 4 20.22 (84.02) 32(8)/1 4 250.22 (84.02) 32(8)/1 4 250.22 (84.02)



Figure 5. Accuracy comparison with respect to the number of parameters on (a) CIFAR100 and (b) Sequential CIFAR100 datasets. Spike-driven Transformer [53] is used as the baseline of spatial-only architecture.

and 83.15% on N-Caltech101, improving the baseline by 1.1% and 1.35% respectively. The enhancement extends to other architectures as well, with SpikingReformer achieving 80.60% on CIFAR10-DVS and 81.95% on N-Caltech101, and QKFormer reaching 83.90% and 84.25% respectively, which is the state-of-the-art result.

5.3. Memory and Energy Analysis

We analyze memory and energy consumption during ImageNet inference. As shown in Table 6, traditional transformer, ViT-B/16, with 32-bit precision requires high memory (351.8 MB) and energy consumption (254.84 mJ). Even with 8-bit precision, Liu [34] still consumes 110.8 mJ of energy. In contrast, spike-based transformers with 32-bit weights and binary activations demonstrate significant energy efficiency. Notably, STAtten integrated with SDT and SDT-V2 achieves similar energy consumption to the baseline architectures, showing that our spatial-temporal processing can be realized without additional energy overhead. Correspondingly, STAtten attains comparable performance with fewer parameters, as detailed in the next section. The methodology for energy calculations is provided in the Supplementary Material.

5.4. Model Capacity

To further evaluate the capacity of our STAtten in spikebased transformers, we compare the accuracy based on the number of trainable parameters through ablation studies. Fig. 5 shows the performance scaling across different model sizes on both (a) standard CIFAR100 and (b) sequential CIFAR100 datasets. Spike-driven Transformer [53] is used as the baseline of spatial-only architecture. As shown in Fig. 5(a), on CIFAR100, STAtten demonstrates strong performance across various model sizes ranging from 2.56M to 22.97M parameters. The accuracy improvement (~ 0.5 -1.0%) is consistently maintained across different architectural scales. For Sequential CIFAR100 (Fig. 5(b)), where temporal dependencies are more crucial, STAtten shows particularly robust performance improvement (~ 3 -5%) even with smaller architectures.

5.5. Limitation

The hardware implementation of STAtten faces challenges despite its improved accuracy. Full spatial-temporal attention without block partitioning is impractical for neuromorphic deployment, as it requires complete temporal information. Our block-wise approach partially addresses this by processing temporal information in steps, where each step corresponds to one block size. However, even with this optimization, deployment remains challenging on traditional neuromorphic chips like TrueNorth [1] and Loihi [6] that process information step-by-step. A potential solution lies in layer-by-layer neuromorphic chips [22, 54, 56, 63] that support layer-wise processing across timesteps. This architectural shift provides an efficient implementation path for STAtten's block-wise processing. Future integration of parallel LIF neurons [15] could further accelerate computation in layer-by-layer architectures.

6. Conclusion

This paper proposes a block-wise spatial-temporal attention mechanism, STAtten, that enhances spike-based transformers. Through block-wise processing and leveraging non-softmax properties, STAtten captures spatial-temporal information without additional overhead. STAtten can be integrated into existing architectures while preserving their energy efficiency, consistently improving performance across different backbones. Experimental results across both static and neuromorphic datasets validate the effectiveness of our approach.

Acknowledgment

This work was supported in part by CoCoSys, a JUMP2.0 center sponsored by DARPA and SRC, the National Science Foundation (CAREER Award, Grant #2312366, Grant #2318152), the DARPA Young Faculty Award and the DoE MMICC center SEA-CROGS (Award #DE-SC0023198).

References

- [1] Filipp Akopyan, Jun Sawada, Andrew Cassidy, Rodrigo Alvarez-Icaza, John Arthur, Paul Merolla, Nabil Imam, Yutaka Nakamura, Pallab Datta, Gi-Joon Nam, et al. Truenorth: Design and tool flow of a 65 mw 1 million neuron programmable neurosynaptic chip. *IEEE transactions on computer-aided design of integrated circuits and systems*, 34 (10):1537–1557, 2015. 1, 8
- [2] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. Vivit: A video vision transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6836–6846, 2021. 1, 4, 6
- [3] Anthony N Burkitt. A review of the integrate-and-fire neuron model: I. homogeneous synaptic input. *Biological cybernetics*, 95:1–19, 2006. 2
- [4] Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Misha Laskin, Pieter Abbeel, Aravind Srinivas, and Igor Mordatch. Decision transformer: Reinforcement learning via sequence modeling. Advances in neural information processing systems, 34:15084–15097, 2021. 1
- [5] Sayeed Shafayet Chowdhury, Isha Garg, and Kaushik Roy. Spatio-temporal pruning and quantization for low-latency spiking neural networks. In 2021 International Joint Conference on Neural Networks (IJCNN), pages 1–9. IEEE, 2021.
- [6] 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. 1, 8
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In 2009 IEEE conference on computer vision and pattern recognition, pages 248–255. Ieee, 2009. 6
- [8] Shikuang Deng and Shi Gu. Optimal conversion of conventional artificial neural networks to spiking neural networks. arXiv preprint arXiv:2103.00476, 2021. 2
- [9] Shikuang Deng, Yuhang Li, Shanghang Zhang, and Shi Gu. Temporal efficient training of spiking neural network via gradient re-weighting. *arXiv preprint arXiv:2202.11946*, 2022.
 7
- [10] Peter U Diehl et al. Fast-classifying, high-accuracy spiking deep networks through weight and threshold balancing. In 2015 International joint conference on neural networks (IJCNN), pages 1–8. ieee, 2015. 2
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929, 2020. 1, 3, 4, 5, 7, 8
- [12] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer* vision, 88:303–338, 2010. 3

- [13] Wei Fang, Zhaofei Yu, Yanqi Chen, Tiejun Huang, Timothée Masquelier, and Yonghong Tian. Deep residual learning in spiking neural networks. *Advances in Neural Information Processing Systems*, 34:21056–21069, 2021. 5, 7
- [14] Wei Fang, Zhaofei Yu, Yanqi Chen, Timothée Masquelier, Tiejun Huang, and Yonghong Tian. Incorporating learnable membrane time constant to enhance learning of spiking neural networks. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2661–2671, 2021. 2, 5, 6, 7
- [15] Wei Fang, Zhaofei Yu, Zhaokun Zhou, Ding Chen, Yanqi Chen, Zhengyu Ma, Timothée Masquelier, and Yonghong Tian. Parallel spiking neurons with high efficiency and ability to learn long-term dependencies. Advances in Neural Information Processing Systems, 36, 2024. 8
- [16] Marylou Gabrié, Andre Manoel, Clément Luneau, Nicolas Macris, Florent Krzakala, Lenka Zdeborová, et al. Entropy and mutual information in models of deep neural networks. Advances in neural information processing systems, 31, 2018. 3
- [17] Bing Han and Kaushik Roy. Deep spiking neural network: Energy efficiency through time based coding. In *European Conference on Computer Vision*, pages 388–404. Springer, 2020. 2
- [18] 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. 6, 1
- [19] Yangfan Hu, Huajin Tang, and Gang Pan. Spiking deep residual networks. *IEEE Transactions on Neural Networks* and Learning Systems, 34(8):5200–5205, 2021. 7
- [20] Yifan Hu, Lei Deng, Yujie Wu, Man Yao, and Guoqi Li. Advancing spiking neural networks toward deep residual learning. *IEEE Transactions on Neural Networks and Learning Systems*, 2024. 1, 2, 5, 7
- [21] Seijoon Kim, Seongsik Park, Byunggook Na, and Sungroh Yoon. Spiking-yolo: spiking neural network for energyefficient object detection. In *Proceedings of the AAAI conference on artificial intelligence*, pages 11270–11277, 2020. 3
- [22] Sangyeob Kim, Soyeon Kim, Seongyon Hong, Sangjin Kim, Donghyeon Han, and Hoi-Jun Yoo. C-dnn: A 24.5-85.8 tops/w complementary-deep-neural-network processor with heterogeneous cnn/snn core architecture and forwardgradient-based sparsity generation. In 2023 IEEE International Solid-State Circuits Conference (ISSCC), pages 334– 336. IEEE, 2023. 8
- [23] Youngeun Kim, Joshua Chough, and Priyadarshini Panda. Beyond classification: Directly training spiking neural networks for semantic segmentation. *Neuromorphic Computing* and Engineering, 2(4):044015, 2022. 2
- [24] 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. 6
- [25] Youngeun Kim, Yuhang Li, Hyoungseob Park, Yeshwanth Venkatesha, Anna Hambitzer, and Priyadarshini Panda. Ex-

ploring temporal information dynamics in spiking neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 8308–8316, 2023. 3

- [26] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 6
- [27] Donghyun Lee, Ruokai Yin, Youngeun Kim, Abhishek Moitra, Yuhang Li, and Priyadarshini Panda. Tt-snn: tensor train decomposition for efficient spiking neural network training. In 2024 Design, Automation & Test in Europe Conference & Exhibition (DATE), pages 1–6. IEEE, 2024. 7
- [28] FF Li, M Andreeto, M Ranzato, and P Perona. Caltech 101 (1.0)[data set]. caltechdata, 2022. 7
- [29] Hongmin Li, Hanchao Liu, Xiangyang Ji, Guoqi Li, and Luping Shi. Cifar10-dvs: an event-stream dataset for object classification. *Frontiers in neuroscience*, 11:244131, 2017. 6, 7
- [30] Yuhang Li, Shikuang Deng, Xin Dong, Ruihao Gong, and Shi Gu. A free lunch from ann: Towards efficient, accurate spiking neural networks calibration. In *International conference on machine learning*, pages 6316–6325. PMLR, 2021.
 2
- [31] Yuhang Li, Yufei Guo, Shanghang Zhang, Shikuang Deng, Yongqing Hai, and Shi Gu. Differentiable spike: Rethinking gradient-descent for training spiking neural networks. *Advances in Neural Information Processing Systems*, 34: 23426–23439, 2021. 2, 6, 7
- [32] Yuhang Li, Youngeun Kim, Hyoungseob Park, Tamar Geller, and Priyadarshini Panda. Neuromorphic data augmentation for training spiking neural networks. In *European Conference on Computer Vision*, pages 631–649. Springer, 2022. 7,
- [33] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In Proceedings of the IEEE/CVF international conference on computer vision, pages 10012–10022, 2021. 7
- [34] Zhenhua Liu, Yunhe Wang, Kai Han, Wei Zhang, Siwei Ma, and Wen Gao. Post-training quantization for vision transformer. Advances in Neural Information Processing Systems, 34:28092–28103, 2021. 8
- [35] Wolfgang Maass. Networks of spiking neurons: the third generation of neural network models. *Neural networks*, 10 (9):1659–1671, 1997. 1
- [36] Qingyan Meng, Mingqing Xiao, Shen Yan, Yisen Wang, Zhouchen Lin, and Zhi-Quan Luo. Training highperformance low-latency spiking neural networks by differentiation on spike representation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 12444–12453, 2022. 6, 7
- [37] Garrick Orchard, Ajinkya Jayawant, Gregory K Cohen, and Nitish Thakor. Converting static image datasets to spiking neuromorphic datasets using saccades. *Frontiers in neuro-science*, 9:159859, 2015. 6, 7
- [38] Jing Pei, Lei Deng, Sen Song, Mingguo Zhao, Youhui Zhang, Shuang Wu, Guanrui Wang, Zhe Zou, Zhenzhi Wu, Wei He, et al. Towards artificial general intelligence with hybrid tianjic chip architecture. *Nature*, 572(7767):106–111, 2019. 1

- [39] Nitin Rathi and Kaushik Roy. Diet-snn: Direct input encoding with leakage and threshold optimization in deep spiking neural networks. arXiv preprint arXiv:2008.03658, 2020. 6
- [40] Kaushik Roy, Akhilesh Jaiswal, and Priyadarshini Panda. Towards spike-based machine intelligence with neuromorphic computing. *Nature*, 575(7784):607–617, 2019. 1
- [41] Abhronil Sengupta, Yuting Ye, Robert Wang, Chiao Liu, and Kaushik Roy. Going deeper in spiking neural networks: Vgg and residual architectures. *Frontiers in neuroscience*, 13:95, 2019. 1
- [42] Xinyu Shi, Zecheng Hao, and Zhaofei Yu. Spikingresformer: Bridging resnet and vision transformer in spiking neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5610–5619, 2024. 1, 2, 4, 5, 6, 7, 8, 3
- [43] Sumit B Shrestha and Garrick Orchard. Slayer: Spike layer error reassignment in time. Advances in neural information processing systems, 31, 2018. 2
- [44] Qiaoyi Su, Yuhong Chou, Yifan Hu, Jianing Li, Shijie Mei, Ziyang Zhang, and Guoqi Li. Deep directly-trained spiking neural networks for object detection. In *Proceedings* of the IEEE/CVF International Conference on Computer Vision, pages 6555–6565, 2023. 2, 3
- [45] Zhengkun Tian, Jiangyan Yi, Jianhua Tao, Ye Bai, Shuai Zhang, and Zhengqi Wen. Spike-triggered nonautoregressive transformer for end-to-end speech recognition. arXiv preprint arXiv:2005.07903, 2020. 1
- [46] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International conference on machine learning*, pages 10347–10357. PMLR, 2021. 7
- [47] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 1, 3
- [48] Qingyu Wang, Tielin Zhang, Minglun Han, Yi Wang, Duzhen Zhang, and Bo Xu. Complex dynamic neurons improved spiking transformer network for efficient automatic speech recognition. In *Proceedings of the AAAI Conference* on Artificial Intelligence, pages 102–109, 2023. 1
- [49] Yujie Wu, Lei Deng, Guoqi Li, Jun Zhu, Yuan Xie, and Luping Shi. Direct training for spiking neural networks: Faster, larger, better. In *Proceedings of the AAAI conference on artificial intelligence*, pages 1311–1318, 2019. 1, 2, 6
- [50] Shiting Xiao, Yuhang Li, Youngeun Kim, Donghyun Lee, and Priyadarshini Panda. Respike: residual frames-based hybrid spiking neural networks for efficient action recognition. *Neuromorphic Computing and Engineering*, 2025. 1
- [51] Man Yao, Guangshe Zhao, Hengyu Zhang, Yifan Hu, Lei Deng, Yonghong Tian, Bo Xu, and Guoqi Li. Attention spiking neural networks. *IEEE transactions on pattern analysis* and machine intelligence, 2023. 7
- [52] Man Yao, Jiakui Hu, Tianxiang Hu, Yifan Xu, Zhaokun Zhou, Yonghong Tian, Bo Xu, and Guoqi Li. Spike-driven transformer v2: Meta spiking neural network architecture inspiring the design of next-generation neuromorphic chips. arXiv preprint arXiv:2404.03663, 2024. 1, 2, 4, 5, 7, 8

- [53] Man Yao, Jiakui Hu, Zhaokun Zhou, Li Yuan, Yonghong Tian, Bo Xu, and Guoqi Li. Spike-driven transformer. Advances in Neural Information Processing Systems, 36, 2024. 1, 2, 3, 4, 5, 6, 7, 8
- [54] Ruokai Yin, Abhishek Moitra, Abhiroop Bhattacharjee, Youngeun Kim, and Priyadarshini Panda. Sata: Sparsityaware training accelerator for spiking neural networks. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 42(6):1926–1938, 2022. 8
- [55] Ruokai Yin, Youngeun Kim, Yuhang Li, Abhishek Moitra, Nitin Satpute, Anna Hambitzer, and Priyadarshini Panda. Workload-balanced pruning for sparse spiking neural networks. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 2024. 2
- [56] Ruokai Yin, Youngeun Kim, Di Wu, and Priyadarshini Panda. Loas: Fully temporal-parallel dataflow for dualsparse spiking neural networks. In 2024 57th IEEE/ACM International Symposium on Microarchitecture (MICRO), pages 1107–1121. IEEE, 2024. 8
- [57] Ruokai Yin, Yuhang Li, Abhishek Moitra, and Priyadarshini Panda. Mint: Multiplier-less integer quantization for energy efficient spiking neural networks. In 2024 29th Asia and South Pacific Design Automation Conference (ASP-DAC), pages 830–835. IEEE, 2024. 2
- [58] Jiqing Zhang, Bo Dong, Haiwei Zhang, Jianchuan Ding, Felix Heide, Baocai Yin, and Xin Yang. Spiking transformers for event-based single object tracking. In *Proceedings of* the IEEE/CVF conference on Computer Vision and Pattern Recognition, pages 8801–8810, 2022. 1
- [59] Jiyuan Zhang, Lulu Tang, Zhaofei Yu, Jiwen Lu, and Tiejun Huang. Spike transformer: Monocular depth estimation for spiking camera. In *European Conference on Computer Vi*sion, pages 34–52. Springer, 2022. 1
- [60] 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*, pages 11062–11070, 2021. 1, 6, 7
- [61] Chenlin Zhou, Liutao Yu, Zhaokun Zhou, Han Zhang, Zhengyu Ma, Huihui Zhou, and Yonghong Tian. Spikingformer: Spike-driven residual learning for transformer-based spiking neural network. arXiv preprint arXiv:2304.11954, 2023. 5
- [62] Chenlin Zhou, Han Zhang, Zhaokun Zhou, Liutao Yu, Liwei Huang, Xiaopeng Fan, Li Yuan, Zhengyu Ma, Huihui Zhou, and Yonghong Tian. Qkformer: Hierarchical spiking transformer using qk attention. *arXiv preprint arXiv:2403.16552*, 2024. 1, 2, 4, 5, 6, 7, 8
- [63] PJ Zhou, Q Yu, M Chen, YC Wang, LW Meng, Y Zuo, N Ning, Y Liu, SG Hu, and GC Qiao. A 0.96 pj/sop, 30.23 k-neuron/mm 2 heterogeneous neuromorphic chip with fullerene-like interconnection topology for edge-ai computing. In 2024 IEEE International Symposium on Circuits and Systems (ISCAS), pages 1–5. IEEE, 2024. 8
- [64] Zhaokun Zhou, Yuesheng Zhu, Chao He, Yaowei Wang, Shuicheng Yan, Yonghong Tian, and Li Yuan. Spikformer: When spiking neural network meets transformer. arXiv preprint arXiv:2209.15425, 2022. 1, 2, 3, 4, 5, 6, 7, 8